



ADSTAR Distributed Storage Manager

SH26-4051-00

Using the Apple Macintosh**
Backup-Archive Client**

Version 2

Note!

Before using this information and the product it supports, be sure to read the general information under "Notices" on page vii.

This book is also available in a softcopy form that can be viewed with the IBM BookManager READ licensed program.

First Edition (November 1996)

This edition applies to Version 2 Release 1 of the ADSTAR Distributed Storage Manager, 5648-020, 5622-112, 5697-078, 5763-SV1, 5733-197, 5686-073, 5655-119, 5765-564, 28H2250, 28H2180, 89G1342, and to any subsequent releases until otherwise indicated in new editions or technical newsletters. Make sure you are using the correct edition for the level of the product.

Order publications through your IBM representative or the IBM branch office serving your locality.

You can send us comments about this book electronically:

- IBMLink from U.S.: STARPUBS at SJSVM28
- IBMLink from Canada: STARPUBS at TORIBM
- IBM Mail Exchange: USIB3VVD at IBMMAIL
- Internet: starpubs@sjsvm28.vnet.ibm.com (or starpubs@sjsvm28.vnet.ibm.com)
- Fax (U.S.): 1-800-426-6209

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1993, 1996. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	vii
Preface	ix
What You Should Know Before Reading This Publication	ix
Style Conventions	x
Referenced Publications	x
Translations	xi
Summary of Changes	xi
Version 2, Release 1, July 1995	xi
Version 2, Release 1, March 1996	xii
Version 2, Release 1, November 1996	xii
Reading Syntax Diagrams	xiii
Chapter 1. Introducing the ADSM Backup-Archive Clients	1
Backing Up and Restoring Files	1
Archiving and Retrieving Files	2
Automating ADSM Tasks	3
Setting Processing Options	3
Chapter 2. Getting Started	5
Starting an ADSM GUI Session	5
Starting ADSM the First Time	5
Starting ADSM Automatically	5
Using the Mouse and Keyboard	6
Using the View Menu	7
Displaying Online Help	8
Changing Your Password	8
Hiding ADSM on the Desktop	9
Ending a Session	9
Chapter 3. Backing Up and Restoring Files	11
Do You Want to Back Up or Archive Files?	11
Backing Up Files and Folders	11
Backing Up New and Changed Files	12
Backing Up Specific Files or Folders	20
Restoring Files or Folders	24
Do You Want To Restore an Active or Inactive Backup?	24
Restoring Files Using a File Specification	25
Restoring Files Using a Tree	27
Restoring Files Using a Folder	28
Restoring Files: Advanced Considerations	30

Chapter 4. Archiving and Retrieving Files	35
Archiving Files	35
Archiving Files Using a File Specification	35
Archiving Files Using a Tree	38
Archiving Files: Advanced Considerations	39
Retrieving Archived Files	41
Retrieving Files	41
Retrieving Files: Advanced Considerations	43
Chapter 5. Automating ADSM Tasks	45
Using Central Scheduling	45
Starting the Client Scheduler	45
Displaying Information About Scheduled Work	46
Displaying Information About Completed Work	47
Specifying Scheduling Options	47
Using Storage Management Policies	49
Understanding Storage Management Policies	49
Displaying Information About Management Classes	51
Choosing a Management Class for Files	55
Choosing a Management Class for Folders	56
Binding and Rebinding Management Classes to Files	57
Chapter 6. Setting Processing Options	59
Understanding Processing Options	59
Communication Options	60
Backup and Archive Processing Options	60
Restore and Retrieve Processing Options	61
Scheduling Options	61
Format and Language Options	62
Command Processing Options	62
Authorization Option	63
Error Processing Options	63
Information Availability Option	63
Setting Options in a File	63
Creating Multiple Preferences Files	64
Creating an Include-Exclude List	65
Excluding System Files	65
Including and Excluding Groups of Files	65
Examples Using Wildcards in EXCLUDE and INCLUDE Patterns	66
Processing INCLUDE and EXCLUDE Options	67
Using Management Classes with INCLUDE	69
Options Reference	70
BOOKS	71
CHANGINGRETRIES	73
COMMMETHOD	74
COMPRESSALWAYS	75
COMPRESSION	76
DATEFORMAT	77

DIRMC	78
DOMAIN	79
ERRORLOGNAME	80
ERRORLOGRETENTION	81
EXCLUDE	82
INCLUDE	83
LANGUAGE	85
MATCHALLCHAR	86
MATCHONECHAR	87
MAXCMDRETRIES	88
NODENAME	89
NUMBERFORMAT	90
PASSWORD	91
POSTSCHEDULECMD	92
PRESCHEDULECMD	93
QUERYSCHEDPERIOD	94
QUIET	95
RETRYPERIOD	96
SCHEDCOMPLETEACTION	97
SCHEDLOGNAME	98
SCHEDLOGRETENTION	99
SLOWINCREMENTAL	100
TAPEPROMPT	101
TCPBUFFSIZE	102
TCPPORT	103
TCPSERVERADDRESS	104
TCPWINDOWSIZE	105
TIMEFORMAT	106
VERBOSE	107
Index	117

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe upon any of the intellectual property rights of IBM may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give one any license to these patents. Send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood NY 10594-1907, USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information that has been exchanged, should contact IBM Corporation, Information Enabling Requests, Dept. M13, 5600 Cottle Road, San Jose CA 95193-0001, USA. Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Trademarks

The following terms, denoted by an asterisk (*) in this publication, are trademarks of the IBM Corporation in the United States, other countries, or both:

Advanced Peer-to-Peer Networking	MVS/ESA
ADSTAR	OpenEdition MVS
AIX	Operating System/2
AIXwindows	Operating System/400
AIX/6000	OS/2
Application System/400	OS/2 Warp
APPN	OS/400
AS/400	RISC System/6000
BookManager	Scalable POWERparallel
Database 2	SP2
IBM	VSE/ESA

The following terms, denoted by a double asterisk (**) in this publication, are trademarks of other companies:

AFS	Transarc Corporation
Apple	Apple Computer, Inc.
Attachmate	Attachmate Corp.
Borland	Borland International, Inc.
CompuServe	CompuServe, Inc.
DECstation	Digital Equipment Corp.
DynaText	Electronic Book Technologies, Inc.
EWS-UX/V	NEC Corporation
Extra!	Attachmate Corp.
Hewlett-Packard	Hewlett-Packard Company
HP-UX	Hewlett-Packard Company
Intel	Intel Corp.
IPX/SPX	Novell, Inc.
IRIX	Silicon Graphics, Inc.
Lotus	Lotus Development Corporation
Lotus Notes	Lotus Development Corporation
Macintosh	Apple Computer, Inc.
MacTCP	Apple Computer, Inc.
Microsoft	Microsoft Corp.
Motif	Open Software Foundation, Inc.
NDS	Novell, Inc.
NetWare	Novell, Inc.
NetWare Directory Services	Novell, Inc.
NetWare Loadable Module	Novell, Inc.
NFS	Sun Microsystems, Inc.
NLM	Novell, Inc.
Novell	Novell, Inc.
Open Desktop	The Santa Cruz Operation, Inc.
OpenWindows	Sun Microsystems, Inc.
PC/TCP	FTP Software, Inc.
SCO	The Santa Cruz Operation, Inc.
SINIX	Siemens Nixdorf Information Systems, Inc.
Solaris	Sun Microsystems, Inc.
SPARC	SPARC International, Inc.
Sun	Sun Microsystems, Inc.
Sun Microsystems	Sun Microsystems, Inc.
SunOS	Sun Microsystems, Inc.
Sun-3	Sun Microsystems, Inc.
Sun-4	Sun Microsystems, Inc.
ULTRIX	Digital Equipment Corp.
WATCOM	WATCOM Systems, Inc.
Windows	Microsoft Corp.
Windows NT	Microsoft Corp.
X Windows	Massachusetts Institute of Technology
X/Open	X/Open Company Limited

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Preface

ADSTAR Distributed Storage Manager (ADSM) is a client/server program product that provides storage management services to customers in a multivendor computer environment. It includes the following programs, which are available on a variety of platforms:

- A *server program* that allows systems to act as a backup and archive server and a migration server for distributed workstations and file servers. The server program provides hierarchical storage management.
- An *administrative client program* that allows an ADSM administrator to control and monitor server activities, define storage management policies for backup, archive, and space management services, and set up schedules to perform those services at regular intervals.
- A *backup-archive client program* that allows users to back up and archive files from their workstations or file servers to ADSM storage, and restore and retrieve backup versions of files and archived copies of files to their local file systems.
- A *hierarchical storage management (HSM) client program* that provides space management services. The HSM client program migrates eligible files to ADSM storage to maintain specific levels of free space on local file systems, and automatically recalls migrated files when they are accessed. It also allows users to migrate and recall specific files. This client is not available on all ADSM-supported platforms.
- An *application programming interface (API)* that allows you to enhance an existing application with storage management services. When an application is registered as a client node with an ADSM server, users can use the application to back up and archive objects, such as databases, to ADSM storage, and restore and retrieve objects from ADSM storage.

This publication provides information about the ADSM Apple Macintosh backup-archive client. ADSM runs on Apple Macintosh System 6.0.7 or later and Power Macintosh.

What You Should Know Before Reading This Publication

This publication is intended for users who need to back up and archive their workstation files.

We assume that you know how to use your workstation, and that you understand the general operation of the operating system.

We also assume that the ADSM backup-archive client is already installed on your system. See the *Installing the Clients* manual for information on installing your client.

Style Conventions

Throughout the book, we use these style conventions:

Figure 1. Style Conventions

Convention	Indicates
bold	A part of a window, or a file name.
UPPER CASE	An ADSM command or option.
<i>italics</i>	A new term.
monospace text	An example in the text.
semicolon (;)	Separates a sequence of actions. You should do these actions one after the other. For example, "click on the ADSM menu; Backup item; Incremental item" means to select the ADSM menu from the menu bar, then select the Backup item from that menu, and finally, select the Incremental item from the resulting menu.
plus sign (+)	Used between two keys to indicate that the two keys it separates should be pressed simultaneously.

Referenced Publications

The publications referenced in this book are listed in Figure 2.

Figure 2. Referenced Publications

Short Title	Publication Title	Order Number
Installing the Clients	<i>ADSTAR Distributed Storage Manager Installing the Clients</i>	SH26-4049
ADSM Messages	<i>ADSTAR Distributed Storage Manager Mes- sages</i>	SH35-0133

All of the ADSM publications are available in online readable format on the *ADSM Online Product Library* CD-ROM, order number SK2T-1878. The ADSM library is also available in softcopy on the following CD-ROMs:

<i>AIX Base Collection Kit</i>	(order number SK2T-2066)
<i>AS/400 Base Collection Kit</i>	(order number SK2T-2171)
<i>MVS Base Collection Kit</i>	(order number SK2T-0710)
<i>OS/2 Base Collection Kit</i>	(order number SK2T-2176)
<i>VM Base Collection Kit</i>	(order number SK2T-2067)
<i>VSE Base Collection Kit</i>	(order number SK2T-0060)
<i>IBM SystemView for AIX</i>	(order number SK2T-1451)

Translations

Selected ADSM publications have been translated into languages other than American English. Contact your IBM representative for more information about the translated publications and whether these translations are available in your country.

Summary of Changes

Major technical changes in this book are indicated by a vertical bar (|) in the margin.

This book was previously released as *ADSM User's Guide and Reference for Apple Macintosh*, SH35-0119.

Version 2, Release 1, July 1995

These are the changes made for ADSM Version 2 Release 1.

Functional Enhancements

- An ADSM hierarchical storage management (HSM) client, which manages storage space on local file systems, is now available on AIX 3.2.3, 3.2.4, and 3.2.5. HSM ensures that sufficient storage space is available on your workstation for new data, by migrating less frequently used files to storage.
- ADSM backup-archive clients are now available for the following workstations:
 - Siemens Nixdorf Information Systems, Inc., SINIX
 - AT&T Global Information Solutions UNIX
 - Silicon Graphics IRIX
 - NEC Corporation EWS-UX/V
 - IBM OpenEdition MVS
 - Microsoft Windows NT 3.5 workstation and Advanced Server on x86 machines
- The following options have been added, although they are not available on all platforms:
 - BOOKS
 - ERRORLOGRETENTION
 - POSTSCHEDULECMD
 - PRESCHEDULECMD
 - SCHEDCOMPLETEACTION
 - SCHEDLOGRETENTION
- Because the Version 2 clients store files in a different format than the Version 1 clients, ADSM updates the attributes for each active file on the ADSM server during the first incremental backup after you install the Version 2 client.

Changes to this Book

- We have moved installation information from this book to a common installation guide, *ADSTAR Distributed Storage Manager Installing the Clients*.
- We have moved the messages from this book to a common message guide, *ADSM Messages*.

- This book has been reorganized and rewritten to make it easier to use.
- To make it easier for you to use the publications in softcopy, we have added a **View Books** menu item to the graphical user interface (GUI) on many ADSM clients.

Version 2, Release 1, March 1996

These are additional changes made for ADSM Version 2 Release 1.

Functional Enhancements

- Support for Power Macintosh has been added.
- The following options have new default values:
 - TAPEPROMPT
 - TCPBUFFSIZE

Version 2, Release 1, November 1996

These are additional changes made for ADSM Version 2 Release 1.

Functional Enhancements

- The COMPRESSALWAYS option has been added.

Reading Syntax Diagrams

This section describes how to read the syntax diagrams used in this book.

To read a syntax diagram, follow the path of the line. Read from left to right and top to bottom.

- The \blacktriangleright — symbol indicates the beginning of a syntax diagram.
- The — \blacktriangleright symbol at the end of a line indicates the syntax diagram continues on the next line.
- The \blacktriangleright — symbol at the beginning of a line indicates a syntax diagram continues from the previous line.
- The — \blacktriangleleft symbol indicates the end of a syntax diagram.

Syntax items, such as a keyword or variable, can be:

- On the line (required element)
- Above the line (default element)
- Below the line (optional element).

Syntax Diagram Description	Example
<p>Abbreviations:</p> <p>Uppercase letters denote the shortest acceptable truncation. If an item appears entirely in uppercase letters, it cannot be truncated.</p> <p>You can type the item in any combination of uppercase or lowercase letters.</p> <p>In this example, you can enter KEYWO, KEYWORD, or KEYWOrd.</p>	<p>\blacktriangleright—KEYWOrd—\blacktriangleleft</p>
<p>Symbols:</p> <p>Enter these symbols exactly as they appear in the syntax diagram.</p>	<ul style="list-style-type: none"> * Asterisk { } Braces : , = Equal Sign - Hyphen () Parentheses . Space
<p>Variables:</p> <p>Italicized lowercase items (<i>var_name</i>) denote variables.</p> <p>In this example, you can specify a <i>var_name</i> when you enter the KEYWORD command.</p>	<p>\blacktriangleright—KEYWOrd—<i>var_name</i>—\blacktriangleleft</p>

Syntax Diagram Description

Example

Repetition:


An arrow returning to the left means you can repeat the item.



A character within the arrow means you must separate repeated items with that character.



A footnote by the arrow references the number of times you can repeat the item.




Note:

1 Specify *repeat* as many as 5 times.

Required Choices:

When two or more items are in a stack and one of them is on the line, you *must* specify one item.



In this example, you *must* choose A, B, or C.

Optional Choice:

When an item is below the line, that item is optional. In the first example, you can choose A or nothing at all.



When two or more items are in a stack below the line, all of them are optional. In the second example, you can choose A, B, C, or nothing at all.



Defaults:


Defaults are above the line. The default is selected unless you override it. You can override the default by including an option from the stack below the line.



In this example, A is the default. You can override A by choosing B or C. You can also specify the default explicitly.

Repeatable Choices:

A stack of items followed by an arrow returning to the left means you can select more than one item or, in some cases, repeat a single item.




In this example, you can choose any combination of A, B, or C.

Syntax Fragments:

Some diagrams, because of their length, must fragment the syntax. The fragment name appears between vertical bars in the diagram. The expanded fragment appears between vertical bars in the diagram after a heading with the same fragment name.



The fragment name:



Chapter 1. Introducing the ADSM Backup-Archive Clients

The ADSTAR Distributed Storage Manager (ADSM) is a client/server program that helps you protect information on your workstation. Using ADSM, you can maintain backup versions of your workstation files that you can restore quickly and easily if the original files are damaged or lost. You can also archive files that you do not currently need on your workstation or that you want to preserve in their current state, and retrieve them when necessary.

The portion of ADSM that provides these services is the *backup-archive client*, a program that allows you to request backup and archive services from an ADSM server.

From the main window, you use the choices in the **ADSM** menu to do your backup and archive tasks.




Figure 3. The ADSM Main Window

The main window is divided into two parts: **Volumes for Backup/Archive** and **Volumes for Restore/Retrieve**. To perform an action, you must select the volumes you want to use from one of these lists, and then select the action from the **ADSM** menu.

Backing Up and Restoring Files

You request *backup* services when you want ADSM to save copies of files that you can restore if your original files are damaged or lost. These copies are called *backup versions* or *backups*.

There are two types of backup using ADSM: incremental and selective.

- *Incremental backup*—Backs up all files and folders on the volumes you specify that are new or that have changed since the last incremental backup.
- *Selective backup*—Backs up only those files you specify.

Running incremental backups on a regular basis can ensure that you have recent backup versions of your current files available at all times. Selective backup is useful when you want to back up only a specific subset of your files.

When you back up files, ADSM also backs up related file and folder attributes, including AppleShare privileges.

If you accidentally erase a file that you need, or if a file on your workstation is damaged or lost, you can *restore* a backup version of the file from the server.

You can restore a single file, a group of files, or all the files in a folder. You can also choose to restore files to the same folder from which you backed them up or to a different folder. You can restore any files you back up and any other files to which another user grants you access.

For more information about backing up and restoring, see Chapter 3, “Backing Up and Restoring Files” on page 11.

Archiving and Retrieving Files

You request *archive* services when you want to save files at a specific point, or when you want to place copies of files on long-term ADSM storage and erase them from your hard disk. These copies are called *archive copies* or *archives*.

Archives are never replaced with more current versions, but are preserved exactly as you store them.

You can archive a single file, a group of files, or all the files in a folder. When you archive a file, you can enter a description of the file that you can use later to retrieve it. You can return an archived file to the same folder from which you archived it or to a different folder.

If you find you need a copy of an archived file back on your workstation, you can *retrieve* the file from the server.

You can retrieve any files you archive and any other files to which another user grants you access.

For more information about archiving and retrieving, see Chapter 4, “Archiving and Retrieving Files” on page 35.

Automating ADSM Tasks

ADSM provides *central scheduling* to automatically perform ADSM tasks. Your ADSM administrator defines schedules. For example, you might have files that you want backed up daily, weekly, or monthly. After defining the appropriate schedules, your administrator associates your client node with those schedules.

You need to decide which schedule you want to use, and then run the *client scheduler*, which tells the server to include your node when it runs the schedule.

For more information about scheduling, see “Using Central Scheduling” on page 45.

Storage management policies are a set of rules that govern how and when your data is backed up and archived. Your ADSM administrator defines these policies, and you associate them with your data.

Your administrator can coordinate scheduled services with the storage management policies you use for your files. For example, if the storage management policy you use for a group of files specifies that they can be backed up only once a week, your ADSM administrator can schedule automatic backup services to occur for those files on a particular day each week.

For more information about storage management policies, see “Using Storage Management Policies” on page 49.

Setting Processing Options

ADSM includes options that control processing for user sessions. For example, you can tell ADSM what characters to use for wildcards and what format to use for dates.

ADSM uses a Preferences file named **ADSM Preferences** to specify options related to your ADSM session. These options are in effect for an entire ADSM session.

For more information about ADSM options, see Chapter 6, “Setting Processing Options” on page 59.

Chapter 2. Getting Started

This chapter contains instructions for starting and ending an ADSM session.

Starting an ADSM GUI Session

To start ADSM:

1 Locate the ADSM icon and its associated files and bring them into view.

2 Double-click on the ADSM icon.

ADSM searches for the **ADSM Preferences** file in the Preferences folder of the system folder, in the system folder itself, and finally in the folder from which ADSM was started. When the Preferences file is located, ADSM starts with the preferences specified in that file.

Note: You can also start ADSM by dragging and dropping the **ADSM Preferences** file onto the ADSM Backup icon.

If you are using System 6.0.7, you can double-click on the **ADSM Preferences** icon to start ADSM.



Figure 4. The ADSM Icon

Your ADSM administrator has the option to force you to use a password to connect to the ADSM server. ADSM prompts you for the password if one is required. Contact your administrator if you do not know your password.

Starting ADSM the First Time

Before you can use ADSM, you must register your client with the ADSM server. The process of setting up a node name and password is called *registration*. Since ADSM provides *open registration* for Macintosh clients, you are prompted for your node name, password, and contact information when you start ADSM for the first time.

Starting ADSM Automatically

You can start the ADSM client scheduler automatically when you start your machine. If the ADSM administrator has defined schedules for your node, starting the client scheduler allows ADSM to automatically back up your machine (or perform other scheduled actions). See Chapter 5, "Automating ADSM Tasks" on page 45 for more information about the client scheduler.

Using the Mouse and Keyboard

A mouse works with ADSM in the same way it works with other standard applications.

Throughout this publication, the term *click* means to press and release the mouse button. For example, “click on **OK**” means to move the mouse pointer to the button labeled **OK**, and then press and release the mouse button. Releasing the mouse button initiates the action associated with the **OK** button.

The term *double-click* means to press and release the mouse button twice on a specified item. For example, “double-click on XYZ” means to move the mouse pointer to the XYZ item and quickly press and release the mouse button twice. Releasing the mouse button the second time initiates the action associated with the XYZ item.

The term *Command-click* means to click on an item while pressing the Command key. In file and folder lists, this deselects individual items embedded in a group of selected items. For example, if you select eight files, and then Command-click on the second file in the list, only the second file is deselected.

The term *Shift-click* means to click on an item while pressing the Shift key. For file and folder lists, this selects groups of files. For example, if you select the first file in a list and then shift-click on the fifth item, all of the items between these two items are selected.

The term *select* means to choose an item from a list or group of items. Usually, the item is highlighted to show that you have selected it. For example, you select a file from a list by moving the mouse pointer to the file and clicking on it. The file is highlighted to indicate that you selected the file.

The term *drag* means to move the mouse while holding down the mouse button, thus moving the selected object.

The term *drag-and-drop* means dragging (moving) an object on top of another object and releasing the mouse button. This relocates an object (such as dropping a file in the trash) or launches an application (such as dropping a file onto an application icon).

The term *scroll* means to move through a list of items in a window by operating the scroll bars with the mouse. For example, clicking the down arrow in the vertical scroll bar moves the contents of the window up into view, scrolling downward through the contents.

Although it is easiest to use a mouse with the ADSM graphical user interface, you can also use the keyboard, or mix using the mouse with keyboard shortcuts. Activate the menus by holding down the Alt key and pressing the underlined letter of the menu you want. The menus indicate the keyboard shortcuts associated with the various actions.

Using the View Menu

When you select the **View** menu in the menu bar, the following menu appears:




Figure 5. The View menu

The first four menu options control the order in which files are displayed. The four choices are:

Sort by folder

Displays files alphabetically by folder name.

This is the default.

Sort by name

Displays files alphabetically by file name.

Sort by size

Displays files in order of size, the largest first, then in descending order.

Sort by modification date

Displays files in order of modification date, showing the most recent files first.

Only one of the four options can be in effect at one time. Initially, a check mark appears to the left of **Sort by folder**. If you select a different option, the check mark moves to that option.

Note: When you are viewing a list of files, you can also click on the headings above the listing to sort by that particular attribute. The current sort key is shown underlined in the window.

The next two menu options control whether ADSM lists both active and inactive backup versions of files or only active backup versions of files during a restore operation. An *active* file is the most recent backup version of a file that existed on your workstation when you ran your last backup. All other backup versions of that file are called *inactive*. If you delete the file from your workstation, the active version becomes inactive the next time you run an incremental backup.

One of the two options is always in effect. When you start a session, **Show active files only** is automatically selected. In the menu, a check mark appears to the left of this option. If you select **Show active and inactive files**, the check mark shifts to that option.

File details lets you control which file details to display in addition to the file name. Initially, the name, size, and modification date shown for each file. When you choose this option, a window appears with the following choices:

- **Size**
- **Modification Date**
- **Management Class**

The three options act as toggle switches. You can select any combination of attributes to appear in file listings. A check mark appears next to each option in effect. You turn off a selection by selecting it again.

The last option, **Refresh volumes**, updates (refreshes) the Volume Information window to show any changes made to your files since the window was first displayed.

The lists of volumes that appear when you start ADSM are updated automatically whenever a major function is completed (at the end of a backup, for example). However, you still need to refresh the list manually on occasion.

For example, when you insert a diskette, you need to refresh before it appears as a new volume beneath **Volumes for Backup/Archive**.

Displaying Online Help

You can display online help in one of two ways. First, you can click on the **Help** button that appears in many windows. The **Help** button displays online information about the current operation.

Second, you can select the **Help** item on the Apple menu or press its keyboard shortcut, Command+H.

If you select **View Books**, ADSM starts an online book reader program, which you can use to read books in the ADSM library (including this one). You or your administrator must install the reader and books on a disk accessible to your system and set the **BOOKS** option in the **ADSM Preferences** file for this help option to work. See "BOOKS" on page 71 for details on how to set up this option.

Changing Your Password

If you are using passwords, you should change them regularly to protect access to your data.

To change your password, click on the **Utilities** menu; **Change password** item. Fill in your old and new passwords and click on **OK**.

An ADSM password can be up to 64 characters. Valid characters are:

Characters	Description
A–Z	any letter, A through Z, upper or lower case
0–9	any number, 0 through 9
+	plus
.	period
_	underscore
-	hyphen
&	ampersand

A password is not case sensitive.

Hiding ADSM on the Desktop

Typically, the Scheduler runs on your Macintosh continuously. To hide the windows for this application, use the following procedure:

- 1** Bring the scheduler to the foreground by clicking on one of its windows, or clicking on **ADSM Scheduler** in the application menu.
- 2** From the application menu, choose **Hide ADSM Scheduler**.

All the ADSM Scheduler windows disappear.

To see the Scheduler again, choose **ADSM Scheduler** from the application menu. This procedure is explained in the Macintosh Guide, available from the Balloon Help menu of the Finder.

Ending a Session

Stop ADSM by selecting the **File** menu; **Quit** item, or by pressing Command+Q.

Chapter 3. Backing Up and Restoring Files

One of ADSM's main uses is to maintain a set of backup versions of the files on your workstation. This allows you to recover older versions of your files if they are lost or damaged.

This chapter covers the various ways you can back up your files using ADSM, and how you can restore them when necessary. It also covers the difference between backing up and archiving files.

Do You Want to Back Up or Archive Files?

When ADSM backs up or archives a file, it sends a copy of the file and its associated attributes to the ADSM server. Backups and archives, though, have different purposes.

Backups are used to protect against file damage. A sequence of backup versions are kept for each file on your workstation (the number of backup versions is set by your ADSM administrator), and older versions are deleted as newer versions are made.

Archive copies, however, are more permanent. They are used to maintain a file in a particular state indefinitely (although your administrator also sets a limit on how long archives are kept). They are useful if you think you might need to go back to a particular version of your files, or if you want to delete a file from your workstation and still be able to retrieve it if necessary. For example, you might need to save some spreadsheets for tax purposes, but because you are not using them, you do not want them left on your workstation. See Chapter 4, "Archiving and Retrieving Files" on page 35 for more information on archive.

Use backups to protect against unforeseen damage to your files, and archives for maintaining more permanent versions of your files.

Backing Up Files and Folders

There are two types of backup using ADSM: incremental and selective.

An *incremental backup* backs up all files that are new or that have changed since the last incremental backup.

A *selective backup* backs up the specific files you indicate.

When you start an ADSM session and the main window first appears, some volumes may already be selected. If so, the selected volumes are specified in your Preferences file. To change the volumes that are selected when you begin, you must edit the DOMAIN option in your Preferences file. See "DOMAIN" on page 79 for more information.

Your administrator might have set up schedules to back up files on your workstation automatically. See Chapter 5, "Automating ADSM Tasks" on page 45 for information

on checking and running the schedules available to you. The following sections cover how to back up files without using a schedule.

Backing Up New and Changed Files

The most effective way to maintain a complete set of backup versions for your most important files is to run incremental backups regularly. An incremental backup backs up the new or changed files on your workstation.

There are two types of incremental backup: *full* and *partial*. A full incremental is also referred to simply as an *incremental* backup.

Do You Want a Full or Partial Incremental?

If you are a casual user of ADSM, and you do not want to worry about the details, skip this section and perform only full incremental backups of your machine. See “Running a Full Incremental Backup” on page 13 for information on how to perform full incremental backups.

There are two main differences between full and partial incremental backups:

- A full incremental ensures that the backups kept on the ADSM server are managed according to your installation's storage management policies (see “Understanding Storage Management Policies” on page 49). It ensures that your backup storage accurately reflects the state of your workstation storage.

A partial incremental does not ensure that the server reflects the exact state of your workstation. For example, files that would normally be backed up during a full incremental might not be backed up during a partial incremental. Also, old files that should be deleted from the server might not be deleted. See “Understanding Partial Incrementals” on page 13 for an explanation of how ADSM determines what happens during a partial incremental.

- A partial incremental should complete more quickly than a full incremental and require less memory. Thus, if you are constrained for time or machine memory, you might want to perform a partial incremental.

Use partial incremental only if you must meet a limited backup time window, and a full incremental takes longer than the allotted time. Periodically perform full incremental backups to bring the ADSM server in line with your workstation's status. For example, if you have only a limited time during the week to perform backups, but extra time on the weekend, you can use partial backups on the weekdays, and then use full incremental backups on the weekends.

Understanding Full Incrementals: In a full incremental, the client queries the ADSM server so that it knows the exact state of your storage as of your last incremental. During a full incremental, ADSM uses this information to:

- Back up new files
- Back up files whose contents have changed
- Back up files whose directory contents have changed

- Expire backup versions on the server for files that have been deleted from the workstation.
- Rebind backup versions to management classes if you have changed management class assignments, even if the associated file is not backed up.

All of this processing is required to maintain your backups according to the policies defined in the management classes.

Understanding Partial Incrementals: In a partial incremental, however, only certain information is considered when deciding whether to back up a file. This is the available type of partial incremental:

- *Incremental by Date.* The client only asks the server for the date and time of the last incremental backup. This is the only information used to determine which files get backed up, so any changes that do not change the date and time are not recognized. If a file's last changed date and time is after that of the last backup, the file is backed up. Otherwise, it is not, even if the file is new to the workstation.

Also, in a partial incremental, files are not expired or rebound on the server, and the copy mode frequency attribute of the management classes is ignored.

Running a Full Incremental Backup

To run a full incremental backup:

- 1 Choose which volumes you want to back up from the **Volumes for Backup/Archive** section of the ADSM window by clicking on them.
- 2 Select the **ADSM** menu; **Incremental backup** item (Figure 6).



Figure 6. Selecting Incremental Backup from the ADSM Menu

ADSM starts the incremental backup, and displays the status of the backup in the Backup Status window (see “Understanding the Backup Status Window” on page 14). ADSM tells you when the backup is complete. The first time you do an incremental backup, ADSM has to back up all your files, which can take a long time. Subsequent backups usually do not take as long.

Understanding the Backup Status Window

During any type of backup, ADSM uses the Backup Status window (Figure 7) to show you what it is doing, and to provide you with statistics on various items.

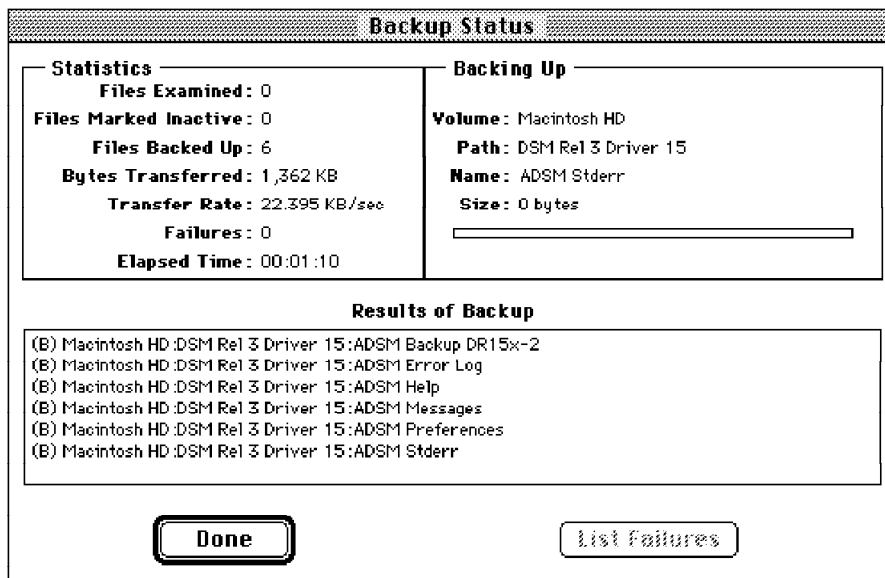


Figure 7. Backup Status Window

The fields in the window are:

Files Examined

The number of files on your workstation that ADSM considered backing up.

Files Marked Inactive

The number of backup versions of files on ADSM storage that were marked inactive because the original files were deleted from your workstation since your last backup.

Files Backed Up

The number of files backed up. This includes folders, which count as two files each.

Bytes Transferred

The number of kilobytes of data sent to the ADSM server.

Transfer Rate

The speed at which bytes were transferred to the ADSM server.

Failures

The number of files ADSM could not back up.

Elapsed Time

The amount of time it took to back up all the files.

Volume

The volume being backed up.

Path

The path within the volume being backed up.

Name

The name of the file being backed up.

Size

The size of the file being backed up.

The lower half of the window, **Results of Backup**, contains a list of the files ADSM processed.

A one-letter code and the path names are shown. The codes are as follows:

(B) Backup

The file was successfully backed up.

(F) Failure

The file was not backed up. If ADSM failed to back up a file, the reason for the failure is shown to the right of the file name. Click on **List Failures** to see a list of the files that ADSM failed to back up.

(I) Inactive

The file was marked inactive on the server because it no longer exists on the workstation.

(R) Rebound

The file was successfully rebound to another management class.

(U) Update

The file's attributes were updated on the server.

Using the Include-Exclude List to Control Processing

Usually, there are files that you do not want to back up. These files might be operating system or application files that could be easily recovered by reinstalling the program, or any other file that you could easily rebuild.

You can use the INCLUDE and EXCLUDE options in the Preferences file to tell ADSM which files to exclude from backup processing. ADSM honors the include-exclude list whether you are doing an incremental or selective backup. ADSM backs up any file that is not explicitly excluded. You do not need to use an INCLUDE option to include specific files unless those files are in a folder you are trying to exclude.

For more information on setting up an include-exclude list, see "Creating an Include-Exclude List" on page 65.

Running an Incremental Backup by Date

To run an incremental backup by date:

- 1** Choose which volumes you want to back up from the **Volumes for Backup/Archive** section of the ADSM window by clicking on them.
Some volumes are already selected. You can deselect these volumes by clicking on them.
- 2** Select the **ADSM** menu; **Incremental by Date** item (Figure 8).



Figure 8. Selecting Incremental by Date from the ADSM Menu

ADSM starts the incremental by date, and displays the status of the backup in the Backup Status window (see “Understanding the Backup Status Window” on page 14).

Incremental Backup: Advanced Considerations

This section discusses some advanced considerations for using incremental backup. You do not need to understand this information in order to use ADSM for basic work.

Understanding How ADSM Uses Volume Labels: ADSM backs up volumes based on the volume label. Each volume label results in ADSM maintaining a separate volume for restore or retrieve in the ADSM Information window. These labels become the names of *restore/retrieve volumes* (or *file spaces*) on the ADSM server.

If you change the label of a volume you have already backed up, ADSM sees it as a new volume and does not relate it to the previous one.

If the volume label does not match a restore/retrieve volume name on the server, ADSM cannot find your files. A mismatch might occur if you relabel your volumes, or if

you access ADSM from a different workstation than the one from which you backed up the files.

Setting Your Default Domain: When you start ADSM, it automatically sets your default domain to the volumes you set with the DOMAIN option in the Preferences file (see “DOMAIN” on page 79). If you do not have the domain option set, the default domain is all local volumes.

Understanding What ADSM Considers a Changed File: For full incrementals, ADSM considers a file to be changed if any of these items are different from the last backup taken:

- File size
- Date or time of last modification
- Icon position within a folder
- Icon color
- Type (document, picture, etc.) or creator
- Finder flags

Simply moving an icon makes the file appear changed. ADSM records the change in position to minimize the problem of multiple icons occupying the same space after the files are restored.

If only the attributes of a file or folder have changed, and not the data, only the attributes are backed up. You may have multiple versions of the same file with the only difference between them being the icon position or color.

If you are running System 7.0 or above, click on **Always snap to grid** in the Views control panel to help control this feature.

Understanding Which Files Are Backed Up: When you request an incremental backup, ADSM backs up a file if all of the following requirements are met:

- The file is not specified on an exclude statement in your include-exclude list. If you do not have an include-exclude list, all files are considered eligible for backup.
- The selected management class contains a backup copy group. See Chapter 5, “Automating ADSM Tasks” on page 45 for more information on management classes and backup copy groups.
- The file meets the serialization requirements defined in the backup copy group. If serialization is *static* or *shared static*, and the file is in use each time ADSM attempts to back it up, the file is not backed up.
- The file meets the mode requirements defined in the backup copy group. If the mode is *modified*, the file must have changed since the last backup. If the mode is *absolute*, the file can be backed up even if it has not changed.
- The file meets the frequency requirements defined in the backup copy group. The minimum number of days since the last backup must elapse before a file is backed up. (Frequency is ignored for partial incrementals.)

When you run an incremental backup, ADSM also backs up all folder information, except for empty folders, which are not backed up. You can exclude the files within a folder, but you cannot exclude a folder from being backed up. Folders are counted in the number of files backed up.

Understanding How Your Files Are Managed: ADSM uses *management classes* to determine how to manage your backups on the ADSM server. Every time you back up a file, the file is assigned a management class. The management class used is either a default chosen for you, or one you assigned to the file using the INCLUDE option in the include-exclude list. The chosen management class must contain a backup copy group in order for the file to be backed up. See Chapter 5, “Automating ADSM Tasks” on page 45 for more information on management classes and how to assign them to files.

AppleShare File Servers: ADSM also backs up AppleShare access privileges. When backing up an AppleShare File Server to ADSM storage, run the backup while logged on as the AppleShare File Server administrator with the “All Privileges” option selected in the administrator’s user window. This ensures access to all files and folders on the File Server volume.

Note: All references to “File Server” in this section refer to the AppleShare File Server, not to the ADSM server. All references to “administrator” in this section refer to the AppleShare File Server administrator, not to the ADSM administrator.

ADSM does not back up the File Server’s System Folder or list of users and groups. You should back up this data onto a 3.5-inch diskette each time it is changed. (This is a good practice whether or not you use ADSM.) Do this from the AppleShare File Server while it is running the AppleShare File Server software:

- 1** Select **Administration** from the Server menu for the server status window.
- 2** Enter your password in the **Admin key** field and click **OK**.
- 3** Select **Save User & Groups As...** from the File menu.
- 4** Insert a 3.5-inch diskette and type a name for the saved list of users and groups. You should use the date as part of the name so you can easily find the most recent version later.

ADSM does not back up the AppleShare File Server system software. You must reinstall it from the original product diskettes if it is damaged or lost; you cannot simply restore it from ADSM. During re-installation, you have the opportunity to load the users and groups from any copies you have on diskette. Once the AppleShare software has been re-installed, you can use ADSM to restore the rest of your data.

System 7 File Sharing Access Privileges: ADSM also backs up System 7 File Sharing Access Privileges, but only if *all* backups are done with File Sharing turned on. Also, to restore folders with their previous access privileges, you must do the restore while File Sharing is on.

Volume Naming Precautions: Unlike other applications on the Macintosh, ADSM is not always case sensitive—it does not see uppercase and lowercase volume names differently. For example, two volumes called “La Pomme” and “la pomme” are seen as identical. Once a volume is backed up to the server, if the case of the volume name changes, ADSM updates the volume name on the server to reflect the change. Keep this in mind when you are naming volumes. Take steps to avoid the possibility of unintended duplicate volumes.

If you do have duplicate volumes of the same name on your desktop, intentionally or otherwise, ADSM cannot determine which one corresponds to the volume of the same name on the server. In this case, ADSM presents an error dialog box if two volumes of the same name are mounted. Remove or rename duplicates as needed.

Backing Up Opened Files: Some files on your system may be in use when you try backing them up (these are called *open files*, because they are opened by an application for its use). Because an open file may change, it is likely that any backup of an open file will not correctly reflect the contents of the file at a given time.

If you find that certain files are not getting backed up because they are open and get changed when you try backing them up (or when the backup schedule starts), you need to consider your options:

- If the file is unimportant, or can be easily rebuilt (for example, a temporary file), you might not care that the file is not backed up.
- If the file is important:
 - You can ensure the file is closed before ADSM tries to back it up. If backups are run according to a schedule, use the PRESCHEDULECMD option to issue a command that closes the file. For example, if the open file is a database, use the database’s quiesce command to shut down the database. You can use the POSTSCHEDULECMD option to restart the application that uses the file after the backup completes.

If you are not using a schedule for the backup, ensure that you close the application that uses the file before you start the backup.

- You can have ADSM back up the file even if it is open and gets changed during the backup. This is only useful if the file will be usable even if it changes during backup. To have ADSM back up these files, assign the files a management class with the serialization **dynamic** or **shared dynamic**. See “Choosing a Management Class for Files” on page 55 for information on assigning management classes, and “Displaying Information About Management Classes” on page 51 for information on determining which management classes are available to you.

Backing Up Specific Files or Folders

You can use ADSM to back up specific files, a group of files with similar names, or entire folders. This is called a *selective backup*.

You can choose the files you want to back up either by using a file specification (the path and name of the file) or by choosing the files or folders from a list.

Backing Up Specific Files Using File Specifications

To back up selected files using file specifications:

- 1 Choose which volumes you want to back up from the **Volumes for Backup/Archive** section of the ADSM window by clicking on them.
- 2 Click on the **ADSM** menu; **Backup by file specification** item (Figure 9).



Figure 9. Selecting Backup by File Specification from the ADSM Menu

- 3 Fill in the name of the file you want to back up.

If the file specification contains folder names, it must contain the complete path minus the volume name. The file specification must end with a file name, not a folder name. Use colons to separate folders.

You can use the standard wildcard characters in the file name portion of the file specification to back up a set of similarly named files (wildcard characters are explained in “Using Wildcard Characters” on page 23). For example, Figure 10 on page 21 shows how to list all files that end in **doc**.

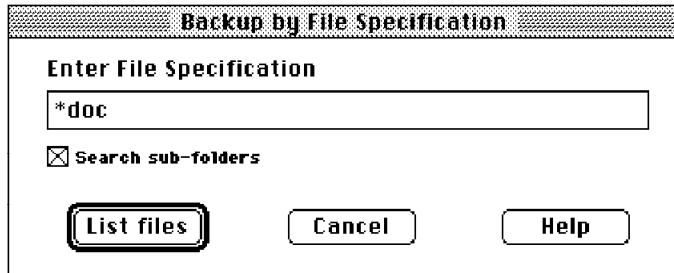


Figure 10. Entering a File Specification for Backup

- 4 Click on **Search sub-folders** to make ADSM search folders within the folder you specify in your path.
- 5 Click on **List files** to see the list of files that match your specification (Figure 11). Select the files you want from this list. (Click on **Select All** to select all the files).

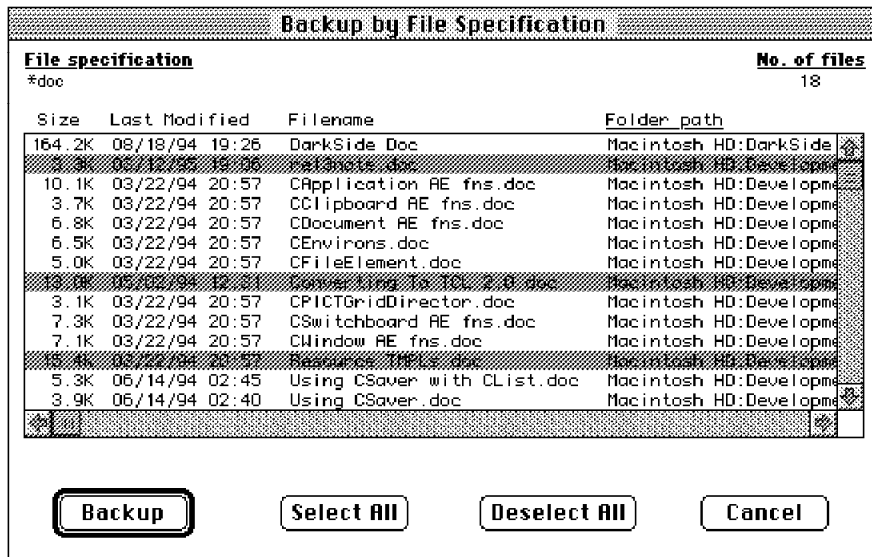


Figure 11. Choosing from the File Specification for Backup List

- 6 Click on **Backup** to back up the selected files. The Backup Status window displays ADSM's processing. (See "Understanding the Backup Status Window" on page 14 for details.)

Backing Up Specific Files Using Trees

To back up specific files or entire folders using trees:

1 Choose which volumes you want to back up from the **Volumes for Backup/Archive** section of the ADSM window by clicking on them.

2 Click on the **ADSM** menu; **Backup by tree** item (Figure 12).

Note: You can also double-click on the volume you want in the ADSM window to open the Backup by Tree window.



Figure 12. Selecting Backup by Tree from the ADSM Menu

3 Select the files and folders you want to back up.

To select or deselect all the files in a folder, Command-click on the folder name. When all the files in a folder are selected, the letter **F** is displayed to the left of the folder name. When some of the files in a folder are selected, the letter **P** is displayed to the left of the folder name. Selecting a folder does not select the folder's sub-folders.

If ADSM does not let you select a file or folder (displayed in italics), it is because your include-exclude list excludes them from backup.

To see a list of the files you selected, click on **List Selections**.

Figure 13 on page 23 shows an example of full and partial folder selections.

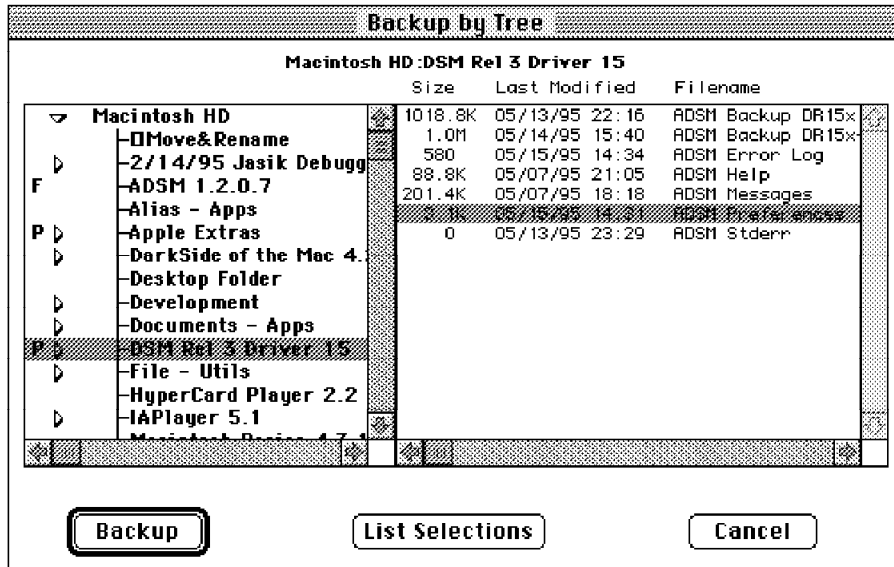


Figure 13. Selecting Files and Folders

- 4 Click on **Backup** to back up the selected files and folders. The Backup Status window displays ADSM's processing. (See "Understanding the Backup Status Window" on page 14 for details.)

Selective Backup: Advanced Considerations

This section discusses some advanced considerations for using selective backup. You do not need to understand this information in order to use ADSM for basic work.

Using Wildcard Characters: You can use the operating system's wildcard characters in file specifications for most ADSM tasks. These characters allow you to select groups of files that have similar names. However, the path itself cannot contain wildcard characters.

The supported wildcard characters are:

- * (Asterisk) This stands for zero or more characters. For example:
 - ***doc** Matches all files that end in **doc**.
 - test*** Matches all files that start with **test**.
- ? (Question mark) This stands for one character. For example:
 - ??**doc** Matches all files that have five characters, ending with **doc**.
 - test?** Matches all files that have five characters, starting with **test**. It would match **test1** and **test2**, but not **test10**.

You can combine the asterisk and question mark in a file specification. For example, **??doc*** matches all files that contain **doc** in the third, fourth, and fifth places. It would match **Exdoc Data** or **Exdoc**, but not **Testdoc**.

Sorting File Lists: You can sort lists of files in ADSM using the items in the **View** menu. You can sort any list of files, whether you are backing up, restoring, archiving, or retrieving them.

Sorting files by last modification date is especially useful if you are trying to restore files to their state as of a particular date. See “Using the View Menu” on page 7 for more information.

Understanding Which Files Are Backed Up: When you request a selective backup, ADSM backs up a file if all of the following requirements are met:

- The file is not specified on an exclude statement in your include-exclude list. If you do not have an include-exclude list, all files are considered for backup.
- The selected management class contains a backup copy group.
- The file meets the serialization requirements defined in the backup copy group. If serialization is *static* or *shared static*, and the file is in use each time ADSM attempts to back it up, the file is not backed up.

When you run a selective backup, ADSM also backs up all folder information. You can exclude the files within a folder, but you cannot exclude a folder from being backed up. Folders are counted in the number of files backed up.

Restoring Files or Folders

You can use ADSM to restore backup versions of specific files or a group of files with similar names, or entire folders.

You can choose the files you want to restore by any of the following methods:

- Use a file specification (the path and name of the file)
- Choose the files or folders from a list
- Use a path to restore a folder and its sub-folders

Do You Want To Restore an Active or Inactive Backup?

Your administrator determines how many backup versions ADSM maintains for each file on your workstation. Frequently, the administrator has ADSM maintain more than one version of each file. Having multiple versions of a file allows you to restore older versions in case the most recent backup is damaged.

ADSM considers the most recent backup version to be the *active* version. Any other backup version is considered an *inactive* version.

Every time ADSM backs up your files, it marks the new backup version as the active backup, and it changes what used to be the active backup to an inactive backup.

When the maximum number of inactive versions is reached, ADSM deletes the oldest inactive version.

If you want to restore a backup version that is inactive, you need to tell ADSM to show both active and inactive versions. Do this by clicking on the **View** menu; **Show active and inactive files** item. You can change back to showing only active versions by clicking on the **View** menu; **Show active files only** item.

If you try to restore both an active and inactive version of a file at the same time, only the active version is restored.

Restoring Files Using a File Specification

If you know the name of the file you want to restore, or you know at least part of the name, you can use the name to restore the file. Restoring a file using a file specification is much like backing one up. The windows are similar, except you are choosing files to restore rather than to back up. To restore a file by its file specification:

- 1 Choose the volumes that contain the backups from the **Volumes for Restore/Retrieve** section of the ADSM window by clicking on them.
- 2 Click on the **ADSM** menu; **Restore by file specification** item (Figure 14).



Figure 14. Selecting Restore by File Specification from the ADSM Menu

- 3 Fill in the name of the file you are restoring.
If the file specification contains folder names, it must contain the complete path minus the volume name. The file specification must end with a file name, not a folder name. Use colons to separate folders.

You can use the standard wildcard characters in the file name portion of the file specification to restore a set of similarly named files (wildcard characters are explained in “Using Wildcard Characters” on page 23). The Restore by File Specification window is similar to the one for backup shown in Figure 10 on page 21.

4 Click on **List files** to see the list of files that match your specification (this list is similar to the one for backup shown in Figure 11 on page 21). Select the files you want from this list. (Click on **Select All** to select all the files).

5 Click on **Restore** and fill in the Restore Parameters window (Figure 15).

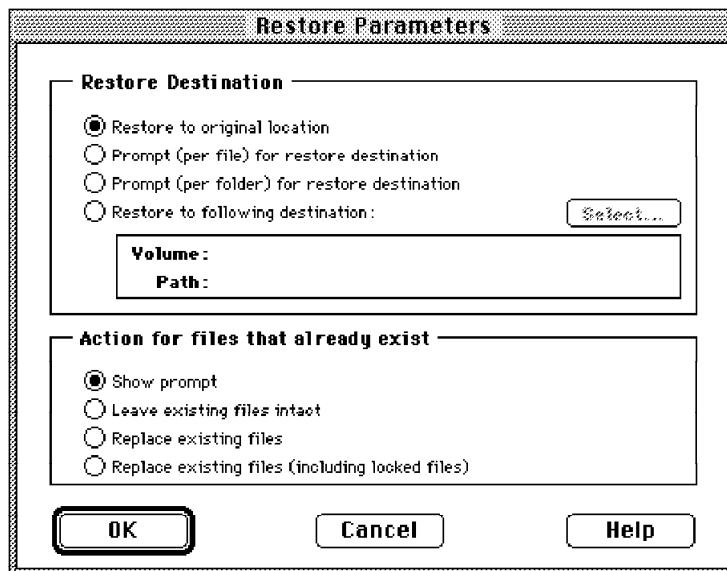


Figure 15. Choosing Restore Parameters

The restore parameters are:

- **Restore Destination.** You can choose where to restore each file:
 - **Restore to original location** restores each file to the same volume and folder from which it was backed up.
 - **Prompt (per file) for restore destination** lets you choose a location for each file as it is restored.
 - **Prompt (per folder) for restore destination** lets you choose a location for each folder of files as the folders are restored.
 - **Restore to following destination** lets you indicate the path to which all files are to be restored.

- **Action for files that already exist.** If you are restoring files that still reside on your disks, you need to tell ADSM how to handle them. You can choose:
 - **Show prompt** causes ADSM to ask you if you want to overwrite the existing file as each file is restored.
 - **Leave existing files intact** causes ADSM to bypass any files that are already on your disk.
 - **Replace existing files** causes ADSM to overwrite existing files without first asking your permission, unless the existing file is read-only.
 - **Replace existing files (including locked files)** causes ADSM to overwrite existing files, including locked files, without first asking your permission.

6 Click on **OK** to restore the files. The Restore Status window (similar to the one for backup; see “Understanding the Backup Status Window” on page 14 for details) displays ADSM's processing.

The lower half of the window, **Results of Restore**, contains a list of the files ADSM processed.

A one-letter code and the path names are shown. The codes are as follows:

(R) Restore

The file was successfully restored.

(F) Failure

The file was not restored. If ADSM failed to restore a file, the reason for the failure is shown to the right of the file name. Click on **List Failures** to see a list of the files that ADSM failed to restore.

Restoring Files Using a Tree

You can use a tree to select the files you want to restore. Restoring a file using a tree is much like backing one up. The windows are the same, except you are choosing files to restore rather than to back up. To restore a file using a tree:

1 Choose the volumes that contain the backups from the **Volumes for Restore/Retrieve** section of the ADSM window by clicking on them.

2 Click on the **ADSM** menu; **Restore by tree** item (Figure 16 on page 28).



Figure 16. Selecting Restore by Tree from the ADSM Menu

- 3** Select the folders and files you want to restore from the Restore by Tree window.
- 4** Click on **Restore** and fill in the Restore Parameters window (Figure 15 on page 26).
- 5** Click on **OK** to restore the files.

Restoring Files Using a Folder

If you want to restore a folder and all of its sub-folders and files, and you know the name of the folder you want to restore, you can use restore by folder.

To restore a folder and its sub-folders:

- 1** Choose the volumes that contain the backups from the **Volumes for Restore/Retrieve** section of the ADSM window by clicking on them.
- 2** Click on the **ADSM** menu; **Restore by folder** item (Figure 17 on page 29).



Figure 17. Selecting Restore by Folder from the ADSM Menu

3 In the section labeled **Source folder**, enter the name of the folder you are trying to restore.

You must enter the complete path, minus the volume name. Use colons to separate folders and sub-folders. For example, to restore the **work** sub-folder in the **project** folder, specify **project:work**.

If you leave the **Path** field blank, ADSM selects the entire volume for restore.

4 In the section labeled **Destination folder**, select the volume and path to which you want the restored files written.

Click on **Select** to view the available folders and select one (Figure 18 on page 30).




Figure 18. Selecting a Destination Folder

You can also create a new target folder by clicking on **New Folder**.

5 In the section labeled **Action for files that already exist**, tell ADSM how to handle files that still reside on your workstation,

6 Click on **OK** to restore the folders.

Restoring Files: Advanced Considerations

This section discusses some advanced considerations for restoring data. You do not need to understand this information in order to use ADSM for basic work.

Authorizing Another User to Restore or Retrieve Your Files

You can authorize a user on another node to restore your backup versions or retrieve your archive copies. This allows you to share files with other people or with other workstations that you use with a different ADSM node name.

To authorize another node to restore or retrieve your files:

1 Click on the **Utilities** menu; **Set authorization** item.

ADSM shows you a list of nodes that you have authorized (see Figure 19 on page 31). The list might be empty.

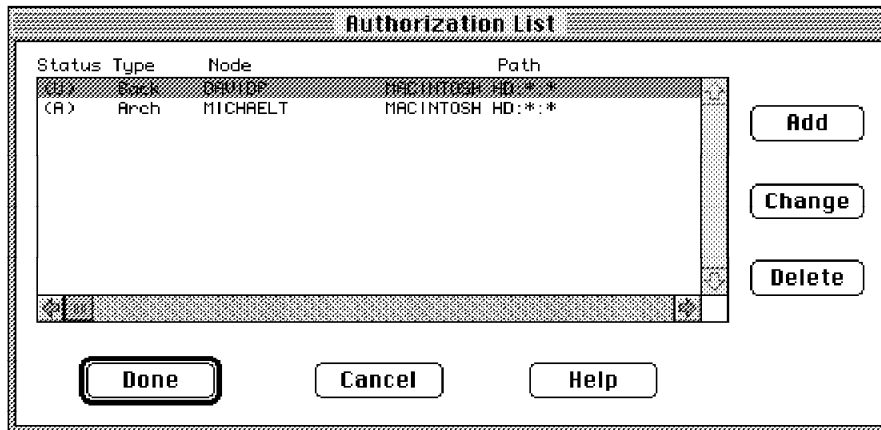


Figure 19. The Authorization List Window

2 Click on **Add** to add a node to the list.

3 Fill in the Authorization window (Figure 20).

Enter the node name and the volume, folder, and file specifications you want the user to have access to. You can give the node access to backups or archives. You must add separate authorizations for backup and archive access, even if you want to give the same node access to the same files for both.

You can authorize all users by using an asterisk (*) for the node name.

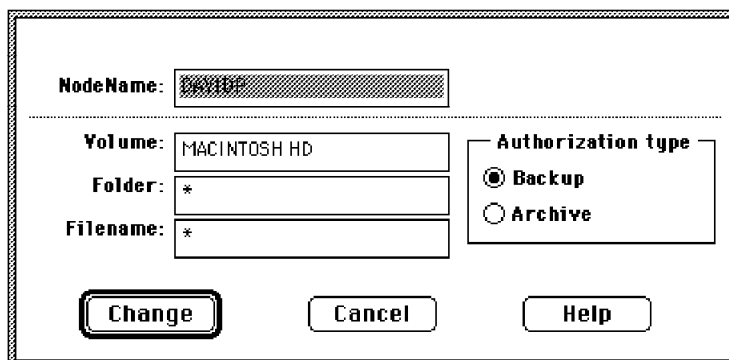


Figure 20. The Authorization Window

4 Click on **Change** in the Authorization window to add the user.

5 Click on **Done** in the Authorization List window to add the node to the authorization list. The node is not added until you do this.

While you are in the Authorization List window, you can add several nodes at once, delete nodes, or change your existing authorizations. All additions, deletions, and changes are processed when you click on **Done**. If you make a mistake, click on **Cancel** to leave the authorization list and start again.

Restoring or Retrieving Another User's Files

After users grant you access to their files on the server, you can restore or retrieve those files by clicking on the **Utilities** menu; **Display volumes for node** item. Enter the name of the node you want to access and click on **OK**.

The other node's restore/retrieve volumes are then displayed in the **Volumes for Restore/Retrieve** section of the ADSM window, and you can restore or retrieve files from them the same way you would your own files.

Restoring or Retrieving Your Files to Another Workstation

ADSM lets you restore or retrieve files you back up from your own workstation when you are using a different workstation. You just need to tell ADSM your node name, because your backups and archives are stored according to your node, not your specific machine. Your ADSM password protects your data.

To restore or retrieve files to another workstation, either change the NODENAME option in your Preferences file to match the current workstation, or simply remove this option from the file. If you remove the NODENAME option, ADSM prompts you for the nodename when it starts. Then enter the node name of the workstation from which you backed up the files.

You can then restore or retrieve files as if you were working on your original workstation.

If you do not want to restore or retrieve the files to the same folder name on the alternate workstation, use a different destination.

You can also have more than one Preferences file with different values for the NODENAME option. See "Creating Multiple Preferences Files" on page 64 for more information.

Restoring a Disk In Case of Disk Loss

ADSM can only recover your files if you can run the ADSM client. If the disk that contains the client is lost (from theft or hardware failure, for example), you must reinstall the client before you can recover your files. If you also lose the disk that contains the operating system and communication software, you must recover them before you can connect to the ADSM server.

To protect yourself against these kinds of losses, you need to put together a set of diskettes that you can use to restore your system to a state that permits you to contact the ADSM server and begin recovering data. These diskettes should contain:

1. A bootable operating system that allows you to perform basic functions.

2. A correctly configured communication program that allows you to establish communications with the ADSM server.
3. An ADSM client with an appropriately customized Preferences file.

Consult your operating system and communication software manuals for help in setting up these diskettes.

Deleting Restore/Retrieve Volumes

If your ADSM administrator gives you backdelete authority, you can delete entire restore/retrieve volumes from ADSM storage. This deletes all the files, both backup versions and archive copies, that are associated with a particular volume.

For example, if you delete the **Macintosh HD** restore/retrieve volume, you are deleting every backup for every file on that volume and every file you archived from that volume. Of course, the original files still exist on your hard disk.

Carefully consider what you are doing before you attempt to delete a restore/retrieve volume.

You cannot delete individual backup versions kept on the ADSM server.

If you decide you do not need a restore/retrieve volume, delete it by clicking on the **Utilities** menu; **Delete restore/retrieve volumes** (Figure 21).



Figure 21. Deleting a File Space

Note: Restore/Retrieve volumes can also be referred to as file spaces.

Overlapping Files

ADSM remembers a file's position on your desktop when it backs up the file. Thus, ADSM tries to restore the file back to its original position. If a file is in the top left-hand corner of your desktop when you back it up, ADSM tries to put it there when you restore it.

But if something else is in that same position when you restore, it is possible that folders and icons could end up overlapping.

If you are having difficulty locating something, click on the **View** menu; **by Name** item in the Finder menu. This lists objects alphabetically, with no overlapping.

Rebuilding the Desktop

If a file you restore does not have the icon it should, you probably need to rebuild the Desktop. Folders with icons pasted onto the Get Info dialog box always require the Desktop to be rebuilt. Do this by restarting your Macintosh with the Option and Command keys held down.

Note: Rebuilding the Desktop loses all **Get Info** comments.

Chapter 4. Archiving and Retrieving Files

Archiving and retrieving files is similar to backing up and restoring them. Many of the windows and concepts are similar. In this chapter, we cover the main archive and retrieve tasks, but where windows and concepts are the same as for backup and restore, we direct you to sections within Chapter 3, "Backing Up and Restoring Files" on page 11.

See "Do You Want to Back Up or Archive Files?" on page 11 for a discussion of the difference between backups and archives.

Archiving Files

To archive files, you need to specifically select the files to archive. You can select the files by using a file specification or by choosing them from a tree.

Your administrator might have set up schedules to automatically archive certain files on your workstation. See Chapter 5, "Automating ADSM Tasks" on page 45 for information on checking and running the schedules available to you. The following sections cover how to archive files without using a schedule.

Archiving Files Using a File Specification

To archive a file using file specifications:

- 1** Choose which volumes contain the files to be archived from the **Volumes for Backup/Archive** section of the ADSM window.
- 2** Click on the **ADSM** menu; **Archive by file specification** item (Figure 22 on page 36).



Figure 22. Selecting Archive by File Specification from the ADSM Menu

- 3** Fill in the name of the file or group of files you want to archive, just as you would do when backing up files.

If the file specification contains folder names, it must contain the complete path minus the volume name. The file specification must end with a file name, not a folder name. Use colons to separate folders.

You can use the standard wildcard characters in the file name portion of the file specification to archive a set of similarly named files (wildcard characters are explained in “Using Wildcard Characters” on page 23). Only the file name portion of the specification may contain wildcard characters.
- 4** Click on the **Search sub-folders** check box, if applicable. This causes ADSM to search folders within the folder you specified in the path.
- 5** Click on **List files** to see the list of files that match your specification. Select the files you want from this list. (Click on **Select All** to select all the files).
- 6** Click on **Archive**.
- 7** Fill in the Archive Options window (Figure 23 on page 37).

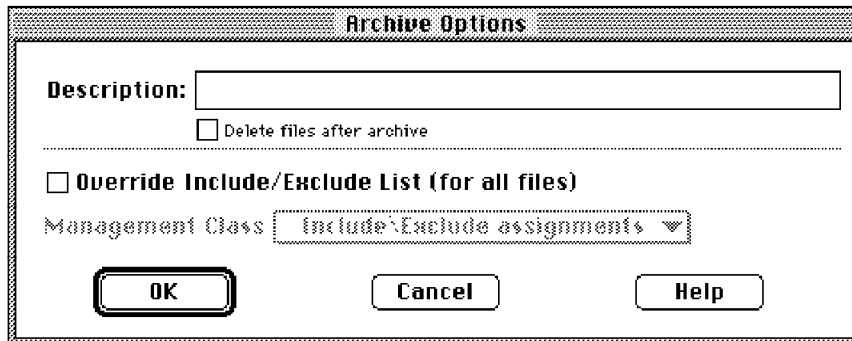


Figure 23. Archive Options Window

The archive options are:

Description

The description you want to keep with the archive copy.

Delete files after archive

ADSM deletes each file from your workstation after it is archived and committed to ADSM server storage.

Management Class

The management class you want assigned to the file (see “Using Storage Management Policies” on page 49 for more information about management classes). You must select **Override Include/Exclude List (for all files)** in order for ADSM to use the management class you choose. Once you select that override option, you can choose a management class by clicking on it. ADSM shows how long archives are kept according to the policy defined for each management class.

- 8 Click on **OK** to start archiving the selected files.

The Archive Status window that is shown is similar to the Backup Status window (see “Understanding the Backup Status Window” on page 14 for details).

A one-letter code and the path names are shown. The codes are as follows:

(A) Archive

The file was successfully archived.

(A/D) Archive/Delete

The file was successfully archived and deleted from your workstation.

(F) Failure

The file was not archived. If ADSM failed to archive a file, the reason for the failure is shown after the code. Click on **List Failures** to see a list of the files that ADSM failed to archive.

- 9 Click on **Done** to return to the ADSM window.

Archiving Files Using a Tree

To archive files assembled from one or more folders:

- 1 Choose which volumes contain the files to be archived from the **Volumes for Backup/Archive** section of the ADSM window.
- 2 Click on the **ADSM** menu; **Archive by tree** item (Figure 24).



Figure 24. Selecting Archive by Tree from the ADSM Menu

- 3 Select the files and folders you want to archive in the same way you would choose them for backup.

To select or deselect all the files in a folder, Command-click on the folder. When all the files in a folder are selected, the letter **F** is displayed to the left of the folder name. When some of the files in a folder are selected, the letter **P** is displayed to the left of the folder name. Selecting a folder does not select the folders it contains.

To see a list of the files you selected, click on **List Selections**.
- 4 Click on **Archive**.
- 5 Fill in the Archive Options window (see Figure 23 on page 37).
- 6 Click on **OK** to start archiving the selected files.

The Archive Status window that is shown is similar to the Backup Status window (see “Understanding the Backup Status Window” on page 14 for details).

7 Click on **Done** to return to the ADSM window.

Archiving Files: Advanced Considerations

This section discusses some advanced considerations for archiving files. You do not need to understand this information in order to use ADSM for basic work.

Deleting Archived Files

You can delete archives if you decide you no longer need them. Unlike backup versions, you can delete individual archives without deleting the entire restore/retrieve volume.

To delete an archive copy:

- 1** Choose which volumes contain the files to be deleted from the **Volumes for Restore/Retrieve** section of the ADSM window.
- 2** Click on the **Utilities** menu; **Delete archived files** item (Figure 25).




Figure 25. Selecting Delete Archived Files from the Utilities Menu

- 3** Fill in the Archive Delete Parameters window (Figure 26 on page 40).

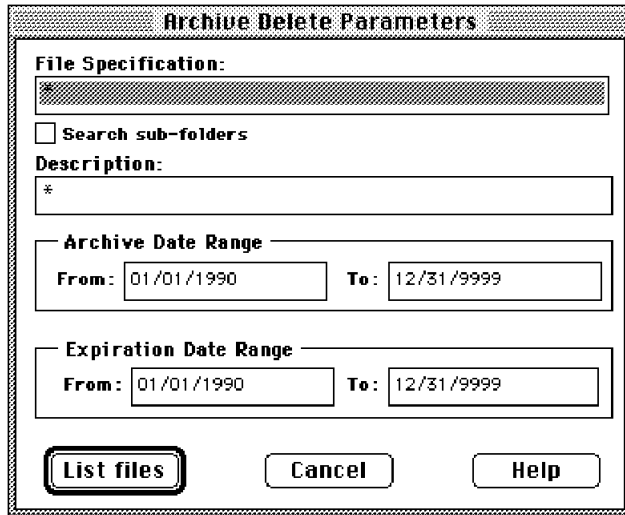


Figure 26. Archive Delete Parameters Window

The Archive Delete Parameters window allows you to build a list of archive copies. You can list all of your archives, or limit the list as much as you want, even to listing the exact file you want to delete, if you know the file's name.

The fields in this window are:

File Specification

The path and name of the file you want to delete. You can use wildcard characters to list similarly named files. Click on **Search sub-folders** to also search folders within the folder you specify in **File Specification**.

Description

The description you gave the file when you archived it. You can use wildcard characters to list files with similar descriptions. The description is case-sensitive.

Archive Date Range

The date range for when you archived the file.

You can limit the list of files by using short date ranges, or, if you cannot remember when you archived the files, you can set a long range to ensure you find the archive files you are looking for.

Expiration Date Range

The date range for when the archive copy is set to expire (set to be deleted automatically on the server).

- 4** Click on **List files** to get a list of the files that meet the parameters you set in the Archive Delete Parameters window.

- 5 Select the files you want to delete in the Archive Delete window (Figure 27 on page 41).

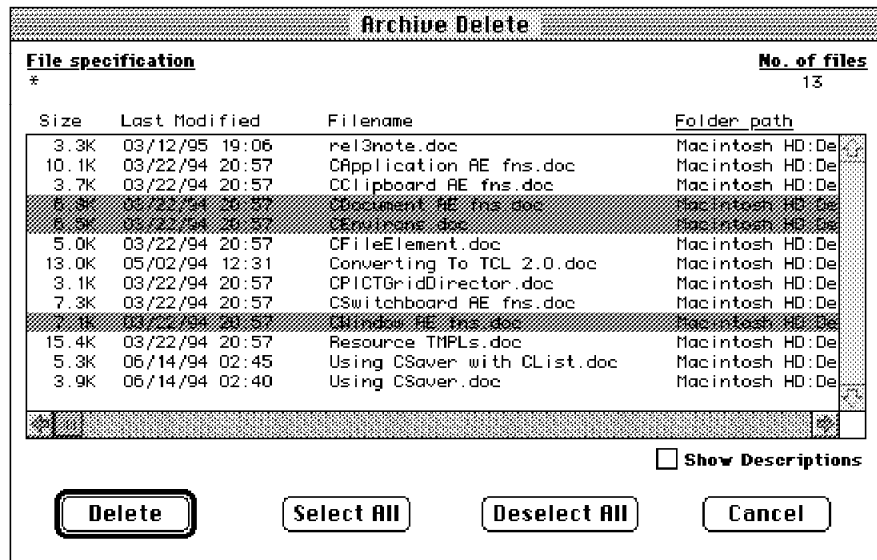


Figure 27. Archive Delete Window

- 6 Click on **Delete** to delete the files.
The Archive Delete Status window shows the results of your action.
A one-letter code and the path names are shown. The codes are as follows:
 - (D) Deleted
The file was successfully deleted.
 - (F) Failure
The file was not deleted.
- 7 Click on **Done** to return to the ADSM window.

Retrieving Archived Files

You *retrieve* a file when you want to return an archive copy from the ADSM server to your workstation.

Retrieving Files

To retrieve an archived file:

- 1** Choose the volumes that contain the archives from the **Volumes for Restore/Retrieve** section of the ADSM window.
- 2** Click on the **ADSM** menu; **Retrieve** item (Figure 28).



Figure 28. Selecting Retrieve from the ADSM Menu

- 3** Fill in the Retrieve Scope window. This window is similar to the Archive Delete Parameters window shown in “Deleting Archived Files” on page 39. You can get a list of archived files using the file name, description, archive date, and expiration date for the file.
- 4** Click on **List files** to get a list of files.
- 5** Select the files you want to retrieve in the Retrieve by File Specification window (similar to Figure 27 on page 41).
- 6** Click on **Retrieve**.
- 7** Fill in the Retrieve Parameters window. This is similar to the Restore Parameters window shown in Figure 15 on page 26. You select where to retrieve the files and what to do if a file is already at the destination.
- 8** Click on **Retrieve** to retrieve the files. The Retrieve Status window shows the results of your action.

A one-letter code and the path names are shown. The codes are as follows:

(R) Retrieved

The file was successfully retrieved.

(F) Failure

The file was not retrieved.

9 Click on **Done** to return to the ADSM window.

Retrieving Files: Advanced Considerations

This section discusses some advanced considerations for retrieving files. You do not need to understand this information in order to use ADSM for basic work.

Understanding the Similarities between Retrieve and Restore

Many of the advanced considerations for retrieving files are the same as they are for restoring files. See these sections for more information:

- “Authorizing Another User to Restore or Retrieve Your Files” on page 30
- “Restoring or Retrieving Another User's Files” on page 32
- “Restoring or Retrieving Your Files to Another Workstation” on page 32

Understanding How Your Archives Are Managed

As with backing up files, ADSM checks the INCLUDE options in your include-exclude list to determine which management class to assign to your archived files. If you do not specifically assign a management class to a file with an INCLUDE option, ADSM assigns the file the default management class.

ADSM can only archive a file if the selected management class contains an archive copy group.

For information on the various management class attributes used to manage your archives, see “Displaying Information About Management Classes” on page 51. For information on assigning management classes, see “Using Management Classes with INCLUDE” on page 69.

Chapter 5. Automating ADSM Tasks

One of the best ways to use ADSM is to automate the tasks you want done on a regular basis. Your ADSM administrator has probably set up some ADSM processes that you can use to do this. For example, ADSM can back up files at the end of each day automatically. Or, it can archive some of your files every Friday.

Central scheduling is a procedure your administrator defines to process ADSM tasks automatically. You start the *client scheduler* on your workstation, which connects you to the server's schedule.

Whereas you can use central scheduling to actually perform your tasks automatically, you can use *storage management policies* to tell ADSM how to manage your backups and archives on the ADSM server. Your ADSM administrator defines these policies, and you associate your data with them. Then, when your data is backed up or archived, it is managed according to the criteria set by the policy.

Using Central Scheduling

An ADSM administrator can set up schedules to perform ADSM tasks automatically. Central scheduling is a cooperative effort between the server and your client node. An ADSM administrator sets up central scheduling on the server and you are responsible for starting the client scheduler on your node. After you start the client scheduler, no further intervention is necessary.

You can also:

- Display information about available schedules
- Display information about work that the schedule has completed
- Modify scheduling options in the Preferences file, if necessary

Starting the Client Scheduler

You must start the client scheduler on your client node to use the server's schedule. To start the Scheduler, double-click on the Scheduler icon. This launches the Scheduler with the options specified in the Preferences file.

Note: You can also drag the Preferences file of your choice and drop it onto the Scheduler icon to start the Scheduler.

When you start the client scheduler, it runs continuously until you close the window by selecting the **File** menu; **Quit** item or by pressing Command+Q.

You can run the scheduler in the background so it does not get interrupted when a pop-up message is displayed. Using a utility, such as the public-domain cron utility, you can launch the scheduler as a background application.

To start the scheduler automatically when the Macintosh is restarted, create an alias for the Scheduler icon and put the alias in the Startup Items folder.

Displaying Information About Scheduled Work

To display information about schedules defined for your client node, click on the **Scheduler** menu; **Query schedules** item. ADSM lists all work scheduled for your client node in the ADSM Scheduler window (Figure 29).

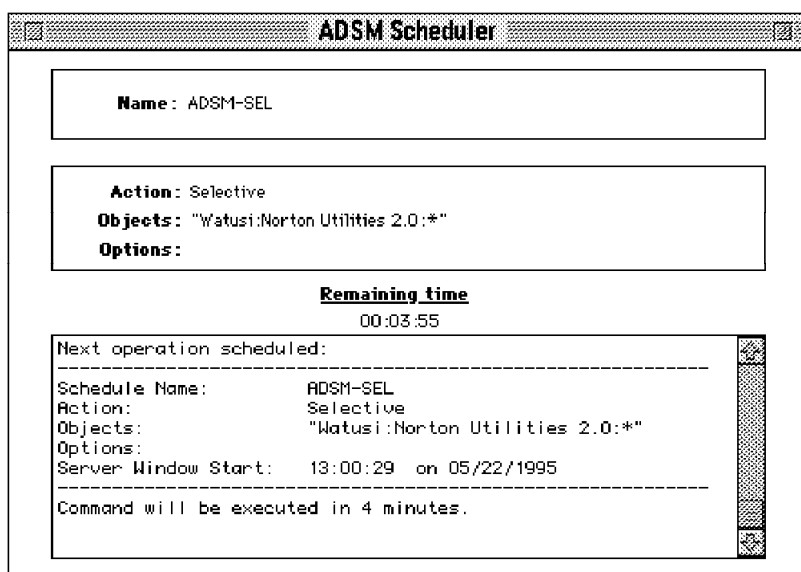


Figure 29. Scheduler Window

The Scheduler window shows the following details about schedules:

Name

The name the administrator has given to the schedule.

Action

The type of task to perform:

- **Incremental**—The schedule incrementally backs up the volumes listed in the **Objects** field.
- **Selective**—The schedule backs up the files listed in the **Objects** field.
- **Archive**—The schedule archives the files listed in the **Objects** field.
- **Command**—The schedule runs the command shown in the **Objects** field.

Objects

The volumes or files to be processed. If volumes are not listed for an incremental backup, ADSM uses the default domain.

Options

Override options that you have set in the Preferences file for scheduling. See "Specifying Scheduling Options" on page 47 for more information about scheduling options.

Displaying Information About Completed Work

When you run the ADSM Scheduler, output from scheduled commands appears in the Scheduler window. Output is also directed to a file called **ADSM Scheduler Log** in the application folder, unless you change the destination or file name using the SCHEDLOGNAME option.

After scheduled work is performed, you can check the schedule log to find out whether all work completed successfully. You can open the file **ADSM Scheduler Log** using TeachText or another Macintosh text editor. This method is preferred, because an editor displays all information, whereas the ADSM Scheduler window truncates the data.

Figure 30 on page 48 shows sample output directed to the schedule log. The output shows the successful completion of a scheduled, selective backup.

The schedule log continues to grow unless you prune it using the SCHEDLOGRETENTION option.

Specifying Scheduling Options

You can modify scheduling options in the Preferences file if you do not like the current values.

For more detailed information about each scheduling option, see Chapter 6, "Setting Processing Options" on page 59.

Changing How the Scheduler Runs

The following options determine how the scheduler runs on your workstation.

If your administrator specifies a value for these options, it overrides what you have in the Preferences file.

QUERYSCHEDPERIOD: The maximum number of hours the client scheduler waits between attempts to contact the server for scheduled work. This option is only valid if you are using the client-polling mode.

In the following example, the client scheduler waits for 5 hours:

```
querschedperiod 5
```

MAXCMDRETRIES: The maximum number of times the client scheduler can retry a schedule that fails.

In the following example, the client scheduler tries 3 times:

```
maxcmdretries 3
```

RETRYPERIOD: The number of minutes the client scheduler must wait between retry attempts for a schedule that fails or between unsuccessful attempts to contact the server.

```

05/18/1995 16:48:46 Querying server for next scheduled event.
05/18/1995 16:48:46 Next operation scheduled:
05/18/1995 16:48:46 -----
05/18/1995 16:48:47 Schedule Name:      ADSM-SEL
05/18/1995 16:48:47 Action:          Selective
05/18/1995 16:48:47 Objects:         "Watusi:Norton Utilities 2.0:*"
05/18/1995 16:48:47 Options:        -schedcompleteaction=shutdown
05/18/1995 16:48:47 Server Window Start: 17:30:29 on 05/18/1995
05/18/1995 16:48:47 -----
05/18/1995 16:48:47 Command will be executed in 2 minutes.
05/18/1995 16:50:47
Executing scheduled command now.
05/18/1995 16:50:48 Selective Backup function invoked.

05/18/1995 16:50:50 Normal File-->      145,461 Watusi:Norton Utilities 2.0:Floppier Sent
05/18/1995 16:50:53 Normal File-->      2,670 Watusi:Norton Utilities 2.0:Icon Sent
05/18/1995 16:50:53 Normal File-->      361,394 Watusi:Norton Utilities 2.0:Norton Backup Sent
05/18/1995 16:50:59 Normal File-->      85,275 Watusi:Norton Utilities 2.0:Norton Encrypt Sent
05/18/1995 16:51:01 Normal File-->      579,644 Watusi:Norton Utilities 2.0:Norton Utilities Sent
05/18/1995 16:51:06 Normal File-->      376,515 Watusi:Norton Utilities 2.0:Norton Utilities Help Sen
05/18/1995 16:51:12 Normal File-->      142,905 Watusi:Norton Utilities 2.0:Speed Disk Sent
05/18/1995 16:51:13 Normal File-->      123,062 Watusi:Norton Utilities 2.0:Wipe Info Sent
05/18/1995 16:51:17 -----
05/18/1995 16:51:17
05/18/1995 16:51:17 Total number of objects inspected:      0
05/18/1995 16:51:17 Total number of objects backed up:          8
05/18/1995 16:51:17 Total number of objects deleted:          0
05/18/1995 16:51:17 Total number of objects failed:          0
05/18/1995 16:51:17 Total number of bytes transferred:      1.7M
05/18/1995 16:51:17 Data transfer time:                  18.43 sec
05/18/1995 16:51:17 Data transfer rate:                   96.23 KB/sec
05/18/1995 16:51:17 Elapsed processing time:              0:00:28
05/18/1995 16:51:17 -----
05/18/1995 16:51:18 ANS6208K Backup completed
05/18/1995 16:51:18 Scheduled event 'ADSM-SEL' completed successfully.
05/18/1995 16:51:18 Sending results for scheduled event 'ADSM-SEL'.
05/18/1995 16:51:19 Results sent to server for scheduled event 'ADSM-SEL'.

```

Figure 30. Sample Schedule Log

In the following example, the client scheduler waits for 15 minutes:

```
retryperiod 15
```

SCHEDLOGNAME: The folder and file name where you want ADSM to store the schedule log. The default is the application folder with a file name of **ADSM Schedule Log**.

In the following example, ADSM puts the schedule log in the **SupDrive** folder with a file name **David Sched Log**:

```
schedlogname "SupDrive:David Sched Log"
```

SCHEDLOGRETENTION: The number of days to keep log file entries from the schedule log before pruning it, and whether to save pruned entries.

In the following example, ADSM keeps entries for 7 days and then prunes them. It discards the pruned entries:

```
schedlogretention 7, D
```

Using Storage Management Policies

Storage management policies are rules that your ADSM policy administrator defines to tell ADSM how to manage backups and archives, based on the needs of users and the business. For example, policies determine:

- Whether a file is eligible for backup or archive services
- How many backup versions to keep
- How long to keep backup versions and archive copies
- Where to place the copies in storage

This section first explains more about storage management policies. Then, it shows you how to display what policies are available to you. Next, you can read about how to choose the best policies for your files and folders. Finally, this section explains some concepts about how ADSM associates your data with policies.

Understanding Storage Management Policies

Storage management policies have several components, as described by the following terms:

Copy group

The specific storage management attributes that describe how the server is to manage backed up or archived data. There are two kinds of copy groups: *backup copy groups* and *archive copy groups*.

Management class

A collection of backup and archive copy groups that establishes similar storage management requirements for backing up and archiving data. Most of the work you do with storage management policies is with management classes.

Policy set

A group of one or more management classes. The ADSM administrator uses a policy set to implement different management classes based on business and user needs.

Policy domain

A group of clients that have similar needs for backup and archive. Policy domains contain one or more policy sets.

Understanding Policy Domains and Sets

An ADSM administrator uses *policy domains* to manage a group of client nodes in a logical way. For example, a domain could consist of the following clients:

- An organization, such as payroll, technical writers, or the marketing team
- A physical location, such as a particular building or floor
- Local area networks, such as all clients associated with a particular file server

ADSM includes a policy domain named STANDARD. Initially, your client node is probably associated with that policy domain. However, your ADSM administrator can define additional policy domains if there are groups of users that have unique backup and archive requirements.

Each policy domain can hold numerous *policy sets*. Only one of these policy sets can be active at a time, called the *active policy set*. Each policy set contains a *default management class* and any number of additional management classes.

Understanding Management Classes and Copy Groups

A *management class* contains the specific requirements for backing up and archiving data. An ADSM administrator can establish separate management classes to meet the backup and archive requirements for different kinds of data, such as:

- System data critical for the business
- Application data that changes frequently
- Report data that management reviews monthly
- Legal information that needs to be retained for a long time and requires a lot of disk space

Each file and folder that you back up and each file that you archive must be associated with a management class. You are responsible for associating files and folders with appropriate management classes. If you do not associate a file with a management class, ADSM uses the default management class in the active policy set. If you do not specify a management class for folders, ADSM uses the management class in the active policy set that specifies the longest retention period.

You associate files with management classes by using an include-exclude list. See “Choosing a Management Class for Files” on page 55 for more information. To associate folders with a management class, you use the DIRMC option. See “Choosing a Management Class for Folders” on page 56 for more information.

Within a management class, the specific backup and archive requirements are in *copy groups*. There are two kinds of copy groups: backup copy groups and archive copy groups. A management class can have one backup copy group, one archive copy group, both, or neither.

Backup Copy Groups: A *backup copy group* contains attributes that ADSM uses during the backup process to determine the following:

- Whether a file is backed up only if it has changed since the last time it was backed up
- How many days must elapse before a file can be backed up again
- How a file is processed during backup if it is in use

It also contains attributes that ADSM uses to manage the backup versions of your files on the server. Those attributes control the following:

- Where the server stores backup versions of your files and folders
- How many backup versions the server keeps of your files and folders
- How long the server keeps backup versions of your files and folders

Archive Copy Groups: An *archive copy group* contains attributes that control the following:

- Whether a file is archived if it is in use
- Where the server stores archived copies of your files
- How long the server keeps archived copies of your files

Displaying Information About Management Classes

Before you choose the management classes you want to use, you need to see which ones are available by clicking on the **Utilities** menu; **Display policy information** item. The information is displayed in the Management Class window (Figure 31).

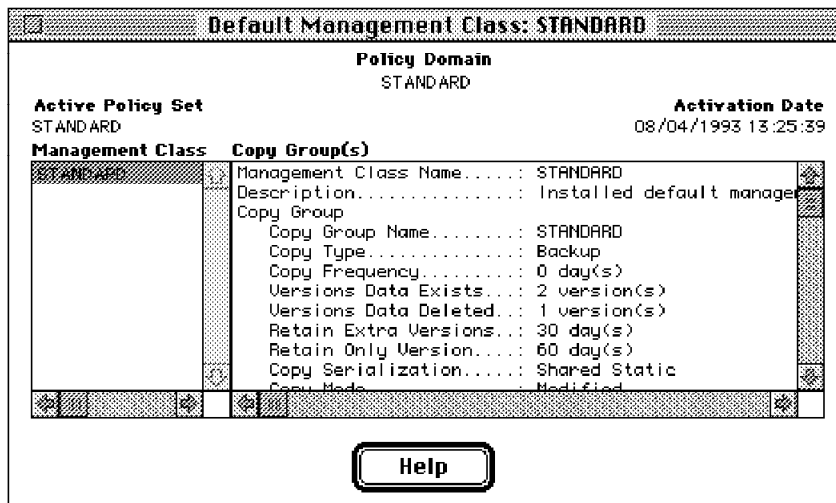


Figure 31. Management Class Window

ADSM displays the following information in this window:

Policy Information

For policy information, you see this information:

Default Management Class

The default management class for the selected policy set is shown in the title bar.

Policy Domain

The name of the policy domain to which this management class belongs.

Active Policy Set

The policy set that is currently active.

Activation Date

The date and time that this policy set became the active one.

Management Class Information

For management class information, you see this information:

Management Class Name

The name of the management class.

Description

A description of the management class.

Copy Groups

You will see information for both backup and archive copy groups:

Copy Group Name

For backup and archive copy groups, the copy group name is always STANDARD.

Copy Type

For backup copy groups, the type is always BACKUP. For archive copy groups, the type is always ARCHIVE.

Copy Frequency

For backup copy groups, copy frequency is the minimum number of days that must elapse between successive incremental backups.

Copy frequency works with the mode parameter that is described later. For example, if frequency is 0, and mode is modified, a file or folder is backed up only if it has been changed since the last incremental backup.

If frequency is 0, and mode is absolute, a file is backed up every time you run an incremental backup against it.

ADSM does not check this attribute for selective backups.

For archive copy groups, the frequency is always CMD (command), that is, there is no restriction on how often you archive a file.

Versions Data Exists

For backup copy groups, the maximum number of different backup versions ADSM retains for files and folders currently on your volume.

If you select a management class that allows more than one backup version, the most current version is called the *active* version. All other versions are called *inactive* versions.

If the maximum number of versions allowed is 5, and you run a backup that creates a sixth version, ADSM deletes the oldest version from server storage.

Versions Data Deleted

For backup copy groups, the maximum number of different backup versions ADSM retains for files and folders that you have erased from your volume. ADSM ignores this parameter as long as the file or folder remains in your volume.

If you erase the file or folder, the next time you run an incremental backup, ADSM changes the active backup version to inactive and erases the oldest versions that are in excess of the number specified by this parameter.

The expiration date for the remaining versions is based on the Retain Extra Versions and Retain Only Version parameters described later.

Retain Extra Versions

For backup copy groups, the number of days ADSM retains all but the most recent backup version. The most recent version is the active version, and active versions are never erased.

If Nolimit is specified, extra backup versions are deleted based on the Versions Data Exists or Versions Data Deleted parameters.

Retain Only Version

For backup copy groups, the number of days ADSM retains the last remaining inactive version of a file or folder.

If Nolimit is specified, the last version is retained indefinitely.

Copy Serialization

For backup and archive copy groups, whether a file can be in use during a backup or archive, and what to do if it is. The value for this parameter can be one of the following:

Static

A file or folder must not be modified during backup or archive. If the object is changed during the backup or archive attempt, it is not backed up or archived.

Shared Static

A file or folder must not be modified during backup or archive. ADSM attempts to perform a backup or archive as many as four additional times, depending on the value specified for CHANGINGRETRIES in your Preferences file. If the object is changed during every backup or archive attempt, it is not backed up or archived.

Dynamic

A file or folder is backed up or archived on the first attempt regardless of whether it changes during backup or archive.

Shared Dynamic

A file or folder is backed up or archived regardless of whether it changes during backup or archive. However, ADSM tries as many as four additional times, depending on the value specified for CHANGINGRETRIES in your Preferences file, to back up or archive the file without the file changing during the attempt. The file is backed up or archived on the last try even if it has changed.

Attention

Be careful about choosing a management class that contains a copy group that specifies shared dynamic or dynamic for serialization.

If you choose a management class that allows a file to be backed up or archived while it is in use, the backup or archive stored on the server might be a fuzzy copy. A fuzzy copy is a backup or archive that does not accurately reflect what is currently in the file. It might contain some, but not all, of the changes. If that is not acceptable, be sure to choose a management class that allows ADSM to create a backup or archive only if the file does not change during backup or archive.

If you restore or retrieve a file that contains a fuzzy copy, the file might not be usable, depending on the file's application.

Copy Mode

For backup copy groups, whether a file or folder is considered for incremental backup only if it has changed since the last backup, or regardless of whether it has changed. ADSM does not check the mode for selective backups. The value for this parameter can be one of the following:

Modified

Indicates that the file is considered for incremental backup only if it has changed since the last backup. A file is considered changed if any of the following are true:

- The date or time of the last modification is different
- The file size is different
- The file attributes, except for the archive attribute, are different

Absolute

Indicates the file is considered for incremental backup regardless of whether it has changed since the last backup.

For archive copy groups, the mode is always absolute, which indicates that a file is archived regardless of whether it has changed since the last archive request.

Copy Destination

For backup and archive copy groups, the name of the destination where ADSM stores the backup or archive. The destination can be either a storage pool of disk devices, or a storage pool of devices that support removable media such as tape.

Retain Versions

For archive copy groups, the number of days an archived file remains in ADSM storage.

When the specified number of days elapse for an archived copy of a file, ADSM deletes it from ADSM storage.

Default Values for Management Class STANDARD

ADSM includes a default management class named STANDARD. Figure 32 shows the default values for the backup and archive copy groups in this management class.

Figure 32. Default Values in the STANDARD Management Class

Attribute	Backup Default	Archive Default
Copy Group Name	STANDARD	STANDARD
Copy Type	BACKUP	ARCHIVE
Copy Frequency	0 days	CMD
Versions Data Exists	2 versions	N/A
Versions Data Deleted	1 version	N/A
Retain Extra Versions	30 days	N/A
Retain Only Version	60 days	N/A
Copy Serialization	shared static	shared static
Copy Mode	modified	absolute
Copy Destination	BACKUPPOOL	ARCHIVEPOOL
Retain Versions	N/A	365 days

Choosing a Management Class for Files

If the default management class meets the backup and archive requirements for all the files on your workstation, you do not need to take any action to associate your files with that management class. ADSM does that automatically when you back up or archive your files.

Using the Include-Exclude List to Assign a Management Class

To assign a management class other than the default to a file or group of files, you use an INCLUDE statement in the include-exclude list. You create the include-exclude list in your Preferences file. For example, to associate all the files in the **Costs** folder with a management class named BUDGET use the following include statement:

```
INCLUDE "...:Costs:*" budget
```

Management class names are not case sensitive.

To specify a management class to be used for all files to which you do not explicitly assign a management class, use the following include statement:

```
INCLUDE * managall
```

where managall is the name of an available management class. For more information about using the include-exclude list, see "Using Management Classes with INCLUDE" on page 69.

Overriding the Management Class for Archived Files

When you archive a file, you can choose to override the assigned management class. ADSM presents a window that allows you to select a different management class. Refer to the **Management Class** field in Figure 23 on page 37 for more information.

Questions to Consider When Choosing a Management Class

If you need to choose a different management class for some of your files, consider the following questions.

- Does the management class contain a backup copy group?
If you attempt to back up a file that is associated with a management class that does not contain a backup copy group, the file is not backed up.
- Does the management class contain an archive copy group?
If you attempt to archive a file that is associated with a management class that does not contain an archive copy group, the file is not archived.
- Does the backup copy group contain attributes that allow your files to be backed up often enough?
Copy mode and copy frequency work together to control how often a file is backed up when you use incremental backup. ADSM does not check those attributes for selective backup.
- Do the copy groups specify either static or shared static for serialization?
If serialization is shared dynamic or dynamic, you might get fuzzy backups or archive copies. Be sure you understand whether that is acceptable.
For example, you might want to use shared dynamic or dynamic for a file to which log records are continuously added. If you used static or shared static, the file might not be backed up at all because it is constantly in use. With shared dynamic or dynamic, the file would be backed up, but the backup version of the file might contain a truncated message.
However, you would not want to use shared dynamic or dynamic for a file if it is critical that the backup version or archive copy contain all changes. That is probably the case for most of your files.
- Does the backup copy group specify an adequate number of backup versions to keep and an adequate length of time to keep them?
- Does the archive copy group specify an adequate length of time to keep archived copies of files?

Choosing a Management Class for Folders

If the management class in your active policy set that contains the longest retention period meets your backup requirements for folders, you do not need to take any action to associate folders with that management class. ADSM does it automatically when it backs up your folders.

If that default management class does not meet your requirements, be sure to choose a management class with an adequate retention period specified for Retain Only Version. You want to be sure that ADSM keeps folders at least as long as it keeps the files associated with them.

To assign a management class other than the default to folders, you use the DIRMC option in your Preferences file. For example, to assign a management class named FOLDER1 to your folders, you would enter:

```
DIRMC FOLDER1
```

See “DIRMC” on page 78 for more information.

Binding and Rebinding Management Classes to Files

When you back up a file for the first time, ADSM binds it to either the default management class or the management class specified for the file in your include-exclude list. *Binding* is the term for associating a file with a management class.

If the backup copy group for the management class instructs ADSM to keep multiple backup versions of the file, and you request multiple backups, the server always has one active backup version (the most current version) and one or more inactive backup versions of the file. All the backup versions of a file are bound to the same management class and are managed based on the attributes in the backup copy group.

When you archive a file, ADSM binds it to the default management class, to the management class specified for the file in your include-exclude list, or to a management class you specify.

There are several instances in which backup versions of a file can be rebound to a different management class. Archived files are never rebound to a different management class. If you change the management class for a file, any previous copies of the file that you have archived remain bound to the management class specified when you archived them.

Rebinding Backup Versions of Files

Backups of files are rebound to a different management class in the following cases. In each case, the files (active and inactive) are not rebound until the next backup.

- You change the management class for the file by specifying a different management class in an INCLUDE statement.

ADSM continues to manage the backups based on the old management class until you run another backup.

- Your ADSM administrator deletes the management class from your active policy set.

ADSM uses the default management class to manage the backup versions when you back up the file again.

- Your ADSM administrator assigns your client node to a different policy domain and the active policy set in that domain does not have a management class with the same name.

ADSM uses the default management class for the new policy domain to manage the backup versions.

Using a Retention Grace Period

ADSM also provides a *backup retention grace period* and an *archive retention grace period* that it uses to help protect your backup and archive data when it is unable to rebind a file to an appropriate management class.

For example, ADSM uses the backup retention grace period in these cases:

- You change the management class for a file, but the new management class does not contain a backup copy group, nor does the default management class.
- The management class to which a file is bound no longer exists, and the default management class does not contain a backup copy group.

ADSM begins using the backup retention grace period when you run an incremental backup.

The backup retention grace period is defined in your policy domain. The default is 30 days. However, your ADSM administrator can choose to change that value to a longer or shorter period.

After ADSM begins managing a file using the backup retention grace period, it does not create any new backup versions of the file. All existing backup versions of the file expire 30 days (or the number of days specified in your policy domain) from the day they are marked inactive.

For archived files, if the management class to which a file is bound no longer exists, and the default management class does not contain an archive copy group, ADSM uses the archive retention grace period defined in your policy domain. The default retention period is 60 days. Your ADSM administrator can choose to change that value to a longer or shorter period.

Chapter 6. Setting Processing Options

You can use defaults for processing options or you can tailor the processing options to meet specific needs. This chapter:

- Provides an overview of processing options
- Explains how to set options in a file

It also includes an options reference section that provides detailed information about each option.

Understanding Processing Options

ADSM processing options reside in the Preferences file. This file, usually called **ADSM Preferences**, contains the following:

- Required communication options and scheduling options
- Options that relate to backup and archive processing
- Options that control processing for sessions with ADSM

If you want to use different options than those specified in the default Preferences file, you can create your own Preferences file.

During the initial installation and setup of ADSM, you are provided with a sample Preferences file called **ADSM Preferences-Sample**. The file is located in the ADSM folder. Make a duplicate of the **ADSM Preferences-Sample** file, and rename the new file **ADSM Preferences** or to a name of your choice. Modify the required options in the new copy. Required options are those that provide information that the backup-archive client program needs to establish communication with an ADSM server.

This section provides an overview of the options you can enter in your Preferences file. The options are arranged in groups by these categories:

- Communication options
- Backup and archive processing options
- Restore and retrieve processing options
- Scheduling options
- Format and language options
- Command processing options
- Authorization options
- Error processing options
- Information availability options

You can use this information to help you decide whether to make any further changes to suit your needs.

Communication Options

You use communication options to specify how your client node communicates with an ADSM server.

ADSM for Apple Macintosh uses the TCP/IP communication protocol. Using TCP/IP requires software called MacTCP (version 2.0.2 or later is recommended). You specify the protocol you want to use with the COMMETHOD option. For more information, see “COMMETHOD” on page 74.

Ask your ADSM administrator for assistance in setting your communication options.

TCP/IP Options

To use the TCP/IP protocol, you must include the TCPSERVERADDRESS option in your Preferences file. The other TCP/IP options have default values. You need to modify these other options only if you want to change the default value.

Figure 33. TCP/IP Options

Option	Description	Reference
TCPBUFFSIZE	Specifies the size, in kilobytes, of the ADSM internal TCP/IP communication buffer.	Page 102
TCPPORT	Specifies the TCP/IP port address for an ADSM server.	Page 103
TCPSERVERADDRESS	Specifies the TCP/IP address for an ADSM server.	Page 104
TCPWINDOWSIZE	Specifies the size, in kilobytes, of the TCP/IP sliding window for your client node.	Page 105

Node Option

Use the following option to specify the client node for which you want to request services

Figure 34. Server and Node Options

Option	Description	Reference
NODENAME	Specifies one of the following: <ul style="list-style-type: none">The name of your workstation so that you can establish communications with the ADSM server.The name of another workstation. Use this if you want to restore or retrieve your files from an ADSM server to a different workstation than the one from which you stored the files.	Page 89

Backup and Archive Processing Options

You can use the following options to control some aspects of backup and archive processing:

Figure 35. Backup and Archive Processing Options

Option	Description	Reference
CHANGINGRETRIES	Specifies the number of retries for attempting to back up or archive a file that is in use.	Page 73
COMPRESSALWAYS	Whether to compress files or send them if compression causes them to grow.	Page 75
COMPRESSION	Specifies whether to compress files before sending them to the ADSM server.	Page 76
DIRMC	Specifies the management class to use for directories. If not specified, ADSM uses the management class with the longest retention period.	Page 78
DOMAIN	Specifies the volumes included in your default client domain for incremental backup.	Page 79
EXCLUDE	Specifies a file or group of files you want to exclude from backup services. Any file in your client domain that is not specifically excluded with this option is considered for backup.	Page 82
INCLUDE	Specifies one of the following: <ul style="list-style-type: none"> Files within a broad group of excluded files that you want to specifically include for backup services. Files to which you want to assign a specific management class and the management class name. A management class to assign to all files to which you do not explicitly assign a management class. 	Page 83
SLOWINCREMENTAL	Specifies a memory-saving backup algorithm for incremental backups for use with the INCREMENTAL command.	Page 100
TAPEPROMPT	Specifies whether you want ADSM to wait for a tape required for a backup or archive to be mounted, or to prompt you for your choice.	Page 101

Restore and Retrieve Processing Options

The following option relates to restore and retrieve processing.

Figure 36. Restore and Retrieve Processing Options

Option	Description	Reference
TAPEPROMPT	Specifies whether you want ADSM to wait for a tape required for a restore or retrieve to be mounted, or to prompt you for your choice.	Page 101

Scheduling Options

Some of the options in the Preferences file regulate central scheduling. You can use these options to provide information that ADSM needs when performing scheduled work. ADSM uses scheduling options only when the Scheduler is running.

Most of the options discussed in this section can be overridden by ADSM server settings after the backup-archive client program on your workstation makes contact with the ADSM server.

Figure 37. Scheduling Options

Option	Description	Reference
MAXCMDRETRIES	Specifies the maximum number of times the client scheduler attempts to process a scheduled command that fails.	Page 88
POSTSCHEDULECMD	Specifies a command to process after running a schedule.	Page 92
PRESCHEDULECMD	Specifies a command to process before running a schedule.	Page 93
QUERYSCHEDULEPERIOD	Specifies the number of hours the client scheduler waits between unsuccessful attempts to contact the ADSM server for scheduled work.	Page 94
RETRYPERIOD	Specifies the number of minutes the client scheduler waits between attempts to process a scheduled command that fails or between unsuccessful attempts to report results to the server.	Page 96
SCHEDCOMPLETEACTION	Specifies an action to take after a schedule has been completed.	Page 97
SCHEDLOGNAME	Specifies the name of the file where schedule log information is stored.	Page 98
SCHEDLOGRETENTION	Specifies the number of days to keep log file entries in the schedule log, and whether to save pruned entries.	Page 99

Format and Language Options

You can use options to choose different formats for date, time, numbers, and different languages if you have the appropriate client installed for that language.

Figure 38. Format and Language Options

Option	Description	Reference
DATEFORMAT	Specifies the format for displaying dates.	Page 77
LANGUAGE	Specifies the language used to display messages.	Page 85
NUMBERFORMAT	Specifies the format for displaying numbers.	Page 90
TIMEFORMAT	Specifies the format for displaying times.	Page 106

Command Processing Options

The following options apply only when you use certain ADSM commands.

Figure 39 (Page 1 of 2). Command Processing Options

Option	Description	Reference
MATCHALLCHAR	Specifies the character to use as a match-all wildcard character.	Page 86

Figure 39 (Page 2 of 2). Command Processing Options

Option	Description	Reference
MATCHONECHAR	Specifies the character to use as a match-one-character wildcard character.	Page 87
QUIET	Specifies that you do not want processing information to appear on your screen.	Page 95
VERBOSE	Specifies that you want processing information to appear on your screen. The alternative is QUIET.	Page 107

Authorization Option

This option controls access to an ADSM server.

Figure 40. Authorization Options

Option	Description	Reference
PASSWORD	Specifies an ADSM password.	Page 91

Error Processing Options

These options control the name of the error log file and how ADSM should treat the entries in the log file.

Figure 41. Error Processing Options

Option	Description	Reference
ERRORLOGNAME	Specifies the path and name of the error log.	Page 80
ERRORLOGRETENTION	Specifies the number of days to keep log file entries in the error log, and whether to save pruned entries.	Page 81

Information Availability Option

This option allows you to make the ADSM online publications available through the ADSM graphical user interface.

If you set the option in the Preferences file, you can read the online publications by clicking on the **Apple** menu; **View Books** item. If you do not set the option, you can still read the online books by starting the online book reader outside of ADSM.

Figure 42. Information Availability Option

Option	Description	Reference
BOOKS	Specifies a command that starts the online book reader.	Page 71

Setting Options in a File

This section provides instructions for setting options in your Preferences file, usually called **ADSM Preferences**.

To set an option in your Preferences file, type the option name, followed by one or more blank spaces, followed by the option parameters. For example:

```
compression  on
nodename     client_a
```

Some options consist of only the option name. For example:

```
verbose
quiet
```

You can write the entire option name or an abbreviation for the option name. For example, for the INCLUDE option, you can write either of the following:

```
include
in
```

The following are additional general rules for entering options in a file:

- You can begin an option in any column of the line.
- You can leave blank lines between options.
- You can enter comments in the Preferences file by entering an asterisk (*) as the first character in a line.
- Put each option on a separate line, and enter all parameters for an option on the same line.

For example, to specify two different volumes as your default client domain, you could enter the following:

```
domain "La Pomme" "Macintosh HD"
```

- Put one or more blank spaces between each parameter.

Creating Multiple Preferences Files

Some tasks require changes to the Preferences file. For example, suppose you wanted to back up your files to one server and archive them to another. Or, suppose you wanted to retrieve files to a different workstation. Instead of editing your Preferences file each time you want to connect to a different server or use a different workstation, you can create multiple Preferences files. Depending on what you want to do, you can launch ADSM with a different Preferences file that suits your needs.

Suppose you have two servers named Alpha and Bravo. Simply create two Preferences files named, for example, **Prefs Alpha** and **Prefs Bravo**. When you want to connect to Bravo, locate **Prefs Bravo** on your desktop and drop it onto the ADSM icon. ADSM starts using all options specified in **Prefs Bravo**, including the server address.

If you are not running System 7.0 or higher, you cannot launch ADSM with drag-and-drop, but you can still create multiple Preferences files and launch ADSM by double-clicking on the Preferences file of your choice.

Note: Remember to make duplicates of your original Preferences file and modify the duplicates. Always keep the original, **ADSM Preferences-Sample**, intact.

Creating an Include-Exclude List

Figure 43 provides brief descriptions of the INCLUDE and EXCLUDE options that you can specify in your Preferences file and references to more detailed information about each option.

Figure 43. Include-Exclude Options

Option	Description	Reference
EXCLUDE	Specifies a file or group of files you want to exclude from backup services. Any file in your client domain that is not specifically excluded with this option is considered for backup.	Page 82
INCLUDE	Specifies one of the following: <ul style="list-style-type: none">Files within a broad group of excluded files that you want to specifically include for backup services.Files to which you want to assign a specific management class and the management class name.A management class to assign to all files to which you do not explicitly assign a management class.	Page 83

Excluding System Files

For Macintosh, we recommend at least the following include-exclude list:

```
exclude "...:Desktop DB"
exclude "...:Desktop DF"
exclude "...:Desktop"
exclude "...:Trash:...:*"
exclude "...:VM Storage"
exclude "...:Norton FileSaver Data"
exclude "...:Norton VolumeSaver Data"
exclude "...:Norton VolumeSaver Index"
```

These files are all system files that can not be recovered without possibly corrupting the operating system.

Note: Any file not explicitly excluded is implicitly included.

Including and Excluding Groups of Files

To specify groups of files that you want to include or exclude, you can use the wildcard characters listed in Figure 44.

Note that this table only applies to include and exclude statements. For information on using wildcard characters in ADSM commands, see "Using Wildcard Characters" on page 23.

Figure 44. Wildcard and Other Special Characters

Character	Function
?	The match-one character matches any single character except the folder separator; it does not match the end of the string. It can only be used in the file name, not in a volume or folder name. (See "MATCHONECHAR" on page 87.)
*	The match-all character matches zero or more characters, but it does not cross a folder boundary. It can only be used in the file name, not in a volume or folder name. (See "MATCHALLCHAR" on page 86.)
...:	Matches all volumes or folders. For example: <p style="margin-left: 40px;">...:.*</p> matches any file in any folder on any volume. The first group specifies the volume and the second group specifies the folder. Note: This is three periods followed by a colon, not an ellipsis followed by a colon.
:	The volume or folder separator. It limits the scope of search for matching all characters and matching all volumes.
[The open character-class character is used to begin the enumeration of a character class. For example: <p style="margin-left: 40px;">xxx[abc] matches xxxa, xxxb, or xxxc.</p>
-	The character-class range includes characters from the first character specified up to and including the last character specified. For example: <p style="margin-left: 40px;">xxx[a-z] matches xxxa, xxxb, xxxc, ... xxxz.</p> Note: For ASCII machines such as the Macintosh, this really means that the characters represented by the hexadecimal codes of 61 (a) through 7A (z) will match. If the last character in the range is less than the first, only the first is used.
]	The close character-class character ends the enumeration of a character class.
\	The literal escape character, when used within a character class, causes the next character to be treated literally. When used outside a character class, it is not specially treated.

Examples Using Wildcards in EXCLUDE and INCLUDE Patterns

Example 1

Exclude all files that end with **doc**, except those found on volume **La Pomme** in the **Documents** folder.

```
exclude "...:.*doc"
include "La Pomme:Documents:.*doc"
```

Example 2

Exclude all files and folders under any **Documents** folder that may exist, *except* for the one file **La Pomme:Documents:Current Resume**. Include this file.


```
exclude "...:Documents:*"
include "La Pomme:Documents:Current Resume"
```

Example 3

Evaluate the following EXCLUDE statements:

```
exclude "Vol[1-4]:...*.cpp"
exclude "Vol2:*.cpp"
exclude "...:Development:*"
```

This example is a bit more complex.

- 1 Vol[1-4]:...*.cpp** matches any **.cpp** file in any folder on the **Vol1**, **Vol2**, **Vol3**, and **Vol4** volumes.
- 2 Vol2:*.cpp** only matches the **.cpp** files found in the root folder of the **Vol2** volume.
- 3 ...:Development:*** matches any file found on any volume that resides under the **Development** folder.

Processing INCLUDE and EXCLUDE Options

ADSM processes the EXCLUDE and INCLUDE options as follows:

- Folders are *not* checked, only files are. Therefore, if you have a volume defined in the DOMAIN option, all folders of that volume are backed up regardless of what you defined for your EXCLUDE and INCLUDE options, even if the files in those folders are not backed up.
- **The order that these options are checked is from the bottom up until a match is found.** If a match is found, the processing stops and looks at whether the option is INCLUDE or EXCLUDE. If the option is INCLUDE, the file is backed up. If the option is EXCLUDE, the file is *not* backed up.
- If a match is *not* found, files listed are implicitly included and therefore backed up.
- If policy administration is in effect, files are backed up according to the default management class or the management class you specify.

Here are some other examples of *bottom up* processing.

Example 1

Assume that **La Pomme** is defined as the domain in the DOMAIN option and you have this defined for the EXCLUDE and INCLUDE options.

```
exclude "...*.cpp"
include "La Pomme:Foo:*.*.cpp"
exclude "La Pomme:Foo:Junk:*.*.cpp"
```

The current file being processed is **La Pomme:Foo:Dev:test.cpp**.

The processing follows these steps:

- 1 Rule 3 (the last INCLUDE or EXCLUDE defined) is checked first because of *bottom up* processing.
- 2 **La Pomme:Foo:Junk:*.cpp** does not match the current file name being processed.
- 3 Processing moves to Rule 2 and checks.
- 4 This time, **La Pomme:Foo:....*.cpp** matches the current file name being processed.
- 5 Processing stops, the option is checked, and it is INCLUDE.
- 6 **La Pomme:Foo:Dev:test.cpp** is backed up.

Example 2

Assume that **La Pomme** is defined as the domain in the DOMAIN option, and you have this defined for the EXCLUDE and INCLUDE options.

```
exclude "...*.cpp"
include "La Pomme:Foo:....*.cpp"
exclude "La Pomme:Foo:Junk:*.cpp"
```

The current file being processed is **La Pomme:Widget:Sample File**.

The processing follows these steps:

- 1 Rule 3 is checked and finds no match.
- 2 Rule 2 is checked and finds no match.
- 3 Rule 1 is checked and finds no match.
- 4 Because no match is found, **La Pomme:Widget:Sample File** falls into the implicit INCLUDE option and is backed up.

Example 3

Assume that **La Pomme** is defined as the domain in the DOMAIN option and you have this defined for the EXCLUDE and INCLUDE options.

```
exclude ".....*.cpp"
include "La Pomme:Foo:....*.cpp"
exclude "La Pomme:Foo:Junk:*.cpp"
```

The current file being processed is **La Pomme:Lib:Src:Module1.cpp**.

The processing follows these steps:

- 1** Rule 3 is checked and finds no match.
- 2** Rule 2 is checked and finds no match.
- 3** Rule 1 is checked and a match is found.
- 4** Processing stops, the option is checked, and it is EXCLUDE.
- 5** **La Pomme:Lib:Src:Module1.cpp** is not backed up.

Using Management Classes with INCLUDE

A management class defines when your files are included in a backup or archive, how long they are kept on the ADSM server, and how many versions of the file the ADSM server should keep.

The ADSM server administrator chooses a default management class. You can override the default management class by specifying your own management class.

This is an example of using a management class:

```
exclude "...:*:Snow"
include "La Pomme:Winter:*:Ice"      mcweekly
include "La Pomme:Winter:December:*Ice" mcdaily
include "La Pomme:Winter:January:*Ice" mcmonthly
include "La Pomme:Winter:Winter:White Snow"
```

The processing follows these steps:

- 1** The file named **White Snow** is backed up because of *bottom up* processing. Because you did not specify a management class, the file is bound to the default management class.
- 2** Any file that ends in **Ice** in the path **La Pomme:Winter:January** is bound to the management class MCMONTHLY.
- 3** Any file that ends in **Ice** in the path **La Pomme:Winter:December** is bound to the management class MCDAILY.
- 4** Any other files that end in **Ice** in any folder under **La Pomme:Winter** are bound to MCWEEKLY.
- 5** Any file that ends in **Snow** (except **La Pomme:Winter:Winter:White Snow**) in any directory is excluded from backup.

To specify your own default management class for files that are not explicitly included, specify:

```
INCLUDE "...:*" mgmt_class_name
```

as the first INCLUDE or EXCLUDE defined.

To specifically define a management class for directories, use the DIRMC option. For details, see “DIRMC” on page 78.

Options Reference

The following sections contain detailed information about each of the ADSM processing options. These options are arranged in alphabetical order. Information for each option includes:

- A description of the option
- A diagram that shows the syntax of the option

The option name is shown in a combination of uppercase and lowercase characters. The uppercase characters indicate the minimum abbreviation you can use for the option name. See “Reading Syntax Diagrams” on page xiii for an explanation of these diagrams.

- Detailed descriptions of the option parameters. If the parameter is a constant, its minimum abbreviation is shown in uppercase letters.
- Example of using the option in the Preferences file

BOOKS

Use the BOOKS option to make the ADSM online publications available through the ADSM graphical user interface.

If you set this option, you can read the online publications by clicking on the **Help** menu; **View Books** item. If you do not set the option, you can still read the online books by starting the online book reader outside of ADSM.

Syntax

▶▶—BOOKS *command*—————▶▶

Parameters

command

The file name and folder of the online book reader. This must be a fully qualified path name, including the volume name. The exact name depends on which reader you are using and where the reader is installed.

The reader available is Electronic Book Technologies DynaText.

The *ADSM Online Product Library* contains the online books and readers for the supported platforms. You or your administrator must install the books and readers separately from the product, as described in the information booklet included with the CD-ROM.

Because of the large amount of disk space required to hold the reader programs and ADSM publications, we recommend that the ADSM administrator install the readers and publications on a LAN server or other machine easily accessible to client users.

Attention

When users select the **View Books** item from the GUI, ADSM starts the online book reader as a separate program.

Online book readers might use a large amount of system resources, because they display graphics and process large files. The reader program might require more resources than are available on your machine, and if run at the same time as ADSM (or any other program), the workstation might freeze.

If your workstation freezes in this manner, do not use **View Books**. Instead, to read the online books, first stop ADSM and other programs that use large amounts of system resources, and then start the program.

You do not have to use the BOOKS option in order to use the ADSM online library.

Also, ADSM does not stop the reader program when the user stops ADSM.

BOOKS

Examples

Preferences file example: books "La Pomme:DynaText Folder:DynaText 2.3"

CHANGINGRETRIES

CHANGINGRETRIES

Use the CHANGINGRETRIES option to specify how many additional times you want ADSM to attempt to back up or archive a file that is in use during the first attempt.

This option relates to serialization, which is an attribute in a management class copy group. ADSM uses this option only when serialization is shared static or shared dynamic.

With shared static, if a file is in use during the first backup or archive attempt, ADSM retries the operation the number of times specified with this option. If the file is in use during each attempt, the operation is not completed.

With shared dynamic, if a file is in use during the first backup or archive attempt, ADSM retries the operation the number of times specified with this option. The backup or archive occurs during the last attempt regardless of whether the file is currently in use.

Syntax

▶▶—Changingretries *numberretries*—————▶▶

Parameters

numberretries

The number of times ADSM retries a backup or archive operation if the file is in use during the first attempt.

The range of values is 0 to 4.

The default is 4.

Examples

Preferences file example: changingretries 3

COMMMETHOD

COMMMETHOD

Use the COMMMETHOD option to set the communication method to Transmission Control Protocol/Internet Protocol (TCP/IP).

Syntax

▶▶—COMMmethod TCPip—————▶▶

Parameters

TCPip

The Transmission Control Protocol/Internet Protocol (TCP/IP) communication method. The following are the communication option settings for the TCP/IP communication method:

COMMmethod	TCPip	
TCPserveraddress		yourserver
TCPport		1500
TCPbuffsize		16

Examples


Preferences file example: comm tcpip

COMPRESSALWAYS

COMPRESSALWAYS

Use the COMPRESSALWAYS option to control what ADSM does when a file grows during compression. You can cause ADSM to either continue compressing or not to send the object if it grows during compression. This option has no effect unless the COMPRESSION option is set to Yes.

Syntax



Parameters

No
Object is not sent if it grows during compression.
This is the default.

Yes
File compression continues even if the file grows as a result of compression.

Examples

Preferences file example: COMPRESSALWAYS Yes

COMPRESSION

COMPRESSION

Use the COMPRESSION option to specify whether ADSM should compress files before sending them to the ADSM server. Compressing your files decreases the amount of data storage required to store backup versions and archived copies of your files. However, it can affect ADSM throughput. Typically, a fast processor on a slow line benefits from compression, but a slow processor on a fast line does not.

This option controls compression only if your ADSM administrator specifies that the choice is determined by your client node.

ADSM does not compress a file if it determines that the compression process is increasing the size of the file. This might happen if a file has already been compressed by another program. These files are sent to the server without being compressed by ADSM.

Syntax

►—COMPression— No Yes—►

Parameters

No

Files are not compressed before being sent to the server. You can also specify No as OFF, False, or 0.

This is the default.

Yes

Files are compressed before being sent to the server. You can also specify Yes as ON, True, or 1.

Examples

Preferences file example: compression yes

DATEFORMAT

DATEFORMAT

Use the DATEFORMAT option to select the format you want ADSM to use for displaying dates.

Syntax

▶▶—DATEformat *format_number*—————▶▶

Parameters

format_number

One of the following date formats:

- 1 Displays the date in this format: MM/DD/YYYY
This is the default.
- 2 Displays the date in this format: DD-MM-YYYY
- 3 Displays the date in this format: YYYY-MM-DD
- 4 Displays the date in this format: DD.MM.YYYY
- 5 Displays the date in this format: YYYY.MM.DD

Examples

Preferences file example: DATEFORM 3

DIRMC

DIRMC

Use the DIRMC option to specify the management class you want ADSM to use. If you do not use this option, ADSM uses the management class in the active policy set of your policy domain that has the longest retention period.

When choosing a management class for individual folders, be sure to choose one that allows ADSM to retain them at least as long as it retains the files associated with them.

Syntax

►—DIRMc *mgmtclassname*—◄

Parameters

mgmtclassname

The name of the management class you want to associate with folders. ADSM uses the management class you specify for all folders it backs up. If you do not use this option, the management class with the longest retention period is associated with folders.

Examples

Preferences file example: DIRM managdir

ERRORLOGNAME

ERRORLOGNAME

Use the ERRORLOGNAME option to specify the file in which you want ADSM to store information about errors that occur during processing.

Syntax

▶—ERRORLOGname *filespec*—————▶◀

Parameters

filespec

The file name and folder where you want ADSM to store error log information.

The default file name is **ADSM Error Log**; it is placed in your ADSM application folder.

Examples

Preferences file example: errorlogname "SupDrive:ADSM Error Log"

ERRORLOGRETENTION


ERRORLOGRETENTION

Use the ERRORLOGRETENTION option to specify:

- The number of days to keep entries in the error log
- Whether to save the pruned entries

The error log is pruned when the first error is written to the log after an ADSM session is started. If the only ADSM session you run is the scheduler, and you run it 24 hours a day, the error log might not be pruned according to your expectations. You must stop the session and restart it to allow the log to be pruned when the next error is written.

Syntax



Parameters

N or *days*

How long to wait before pruning the error log.

N

Do not prune the error log. This allows the error log to grow indefinitely.

This is the default.

days

The number of days to keep log file entries before pruning the log.

The range of values is 0 to 9999.

D or S

Whether to save the pruned entries.

If you specify anything but D or S, ADSM uses the default D. You must use a space or comma to separate this parameter from the previous one.

D

Discard the log entries when pruning the log.

This is the default.

S

Save the log entries when pruning the log.

ADSM copies the entries pruned from the error log to **ADSM Error Pruned Log**.

Examples

Preferences file example: ERRORLOGRETENTION 400 S

EXCLUDE

EXCLUDE

Use the EXCLUDE option to exclude a file or group of files from ADSM backup services. When you back up files, any files you exclude with this option are not considered for backup. For example, you might want to exclude some of the Macintosh Operating System files.

You can use wildcard characters to exclude a broad range of files. See “Including and Excluding Groups of Files” on page 65 for a list of wildcards you can use. Then, if necessary, you can use the INCLUDE option to make exceptions.

See “Creating an Include-Exclude List” on page 65 for more information. See also the list of files that should always be excluded in “Excluding System Files” on page 65.

Syntax

▶▶ `EXclude pattern` ◀◀

Parameters

pattern

The file or group of files that you want to exclude from backup services. The pattern must end with a file specification.

If the pattern contains a blank space, you must enclose the pattern in quotation marks.

Examples

Preferences file example:

```
exclude "...:Desktop DB"  
exclude "...:Norton FileSaver Data"
```


INCLUDE

INCLUDE


Use the INCLUDE option to specify one of the following:

- Files within a broad group of excluded files that you want to specifically include for backup services.
- Files to which you want to assign a specific management class and the management class name.
- A management class to assign to all files to which you do not explicitly assign a management class.

If you do not assign a specific management class to files, ADSM uses the default management class in the active policy set of your policy domain.

See “Creating an Include-Exclude List” on page 65 for more information.

Syntax



Parameters

pattern

The specific files you want to include for backup services, or the file or group of files to which you want to assign a specific management class.

You must end the pattern with a file specification.

If the pattern contains a blank, you must enclose the pattern in quotation marks.

mgmtclassname

The name of the management class you want to assign to the files. If you do not specify a management class, ADSM uses the default management class.

Examples

Preferences file example:

Example 1

Explicitly include a subset of files that are excluded:

```
exclude "...:Test*"
include "...:Data:Test:Test*"
```

Example 2

Assign a management class to a group of files:

```
include "...:System:...:*" mgmtcls1
```

Example 3

Specify a management class to assign to all files to which you have not assigned a specific management class:

INCLUDE

```
include "...:*" managall
```

This must be the *first* statement in your include-exclude list.

Use this method to assign a management class to all files to which you have not already assigned a specific management class and for which you do not want ADSM to use the default management class.

LANGUAGE

LANGUAGE

Use the LANGUAGE option to choose the national language for presenting text.

Syntax

▶▶—LANGUage *language*—————▶▶

Parameters

language

The language to use for displaying text. Currently, only American English (AMENG) can be used, which is the default if you do not use this option.

Examples

Preferences file example: LANGUAGE AMENG

MATCHALLCHAR

MATCHALLCHAR

Use the MATCHALLCHAR option to specify the character to use as a match-all wildcard character. For example, **File*** matches **File1**, **FileA1**, and **File**.

This option must appear before any INCLUDE or EXCLUDE statements in the Preferences file.

Syntax

▶—MATCHAllchar *character*—————▶◀

Parameters

character

A character to be used as a match-all wildcard character.

You can use any valid ASCII character for MATCHALLCHAR except Control+X or Control+Y. Do not use the same character for both MATCHALLCHAR and MATCHONECHAR.

The default is an asterisk (*).

Examples

Preferences file example: matchallchar @

MATCHONECHAR

MATCHONECHAR

Use the MATCHONECHAR option to specify the character to use as a match-one wildcard character. For example, **File?** matches **File1**, but not **File** or **File10**.

This option must appear before any INCLUDE or EXCLUDE statements in the Preferences file.

Syntax

►► MATCHOnechar *character* ◀◀

Parameters

character

The character to be used as a match-one wildcard character.

You can use any valid ASCII character for MATCHONECHAR, except Control+X or Control+Y. Do not use the same character for both MATCHALLCHAR and MATCHONECHAR.

The default is a question mark (?).

Examples

Preferences file example: matchonechar #

MAXCMDRETRIES

MAXCMDRETRIES

Use the MAXCMDRETRIES option to specify the maximum number of times you want the client scheduler on your workstation to attempt to process a scheduled command that fails.

ADSM uses this option only when the Scheduler is running.

Your ADSM administrator can also set this option. If your ADSM administrator specifies a value for this option, that value overrides what you specify in the Preferences file after your client node successfully contacts the ADSM server.

Syntax

▶—MAXCMDRetries *maxcmdretries*—————▶

Parameters

maxcmdretries

The number of times the client scheduler can attempt to process a scheduled command that fails.

The range of values is 0 to 9999.

The default value is 2.

Examples

Preferences file example: MAXCMDR 4

NODENAME

NODENAME

Use the NODENAME option to identify your workstation to the server. You can also use the NODENAME option if you want to restore or retrieve files from the server while you are working at a different workstation. When you use the NODENAME option, ADSM prompts you for the ADSM password assigned to the node you specify if a password is required.

If you do not specify this option in the Preferences file, ADSM prompts you for the nodename when it starts.

Syntax

▶▶—NODename *nodename*—————▶▶

Parameters

nodename

A 1 to 64 character name to identify the node for which you want to request ADSM services.

Examples

Preferences file example: NODENAME cougar

NUMBERFORMAT

NUMBERFORMAT

Use the NUMBERFORMAT option to specify the format you want ADSM to use for displaying numbers.

Syntax

►—NUMBERFORMAT *number*—————►

Parameters

number

One of the following number formats:

- 1 Displays numbers in this format: 1,000.00
This is the default.
- 2 Displays numbers in this format: 1,000,00
- 3 Displays numbers in this format: 1 000,00
- 4 Displays numbers in this format: 1 000.00
- 5 Displays numbers in this format: 1.000,00
- 6 Displays numbers in this format: 1'000,00

Examples

Preferences file example: num 4

PASSWORD

PASSWORD

Use the PASSWORD option to specify an ADSM password.

If this option is not used, you are prompted for a password when you start an ADSM session if your administrator has set authentication to on.

Be careful about using this option. If ADSM prompts you for a password, ADSM ensures the password is not displayed as you type it in. However, the password cannot be hidden on the password option.

Syntax

▶▶—PASsword *password*—————▶▶

Parameters

password

A 1 to 64 character password. Valid characters are:

Characters	Description
A–Z	any letter, A through Z, upper or lower case
0–9	any number, 0 through 9
+	plus
.	period
_	underscore
-	hyphen
&	ampersand

A password is not case sensitive.

Examples

Preferences file example: password secretword

POSTSCHEDULECMD

POSTSCHEDULECMD

Use the POSTSCHEDULECMD option to specify a command that ADSM should process after running a schedule.

If you want ADSM to wait for the command to complete before continuing with other processing, specify POSTschedulecmd. If you don't want it to wait, specify POSTNschedulecmd.

Syntax

► — [POSTschedulecmd] — "cmdstring" —►
 [POSTNschedulecmd]

Parameters

"cmdstring"

The command to process.

You can use only one POSTSCHEDULECMD option. You can also issue a command before a schedule using the PRESCHEDULECMD option.

Use a blank or null string for *cmdstring* if you want to prevent any commands that the ADSM administrator uses for POSTSCHEDULECMD and PRESCHEDULECMD from running. If you specify a blank or null string on *either* option, it prevents the administrator from using a command on *both* options.

However, if your administrator uses a blank or null string on the POSTSCHEDULECMD option, that prevents you from running a post-schedule command.

If the command string contains blanks, you must enclose it in double quotes. If you have double quotes within the command string, use single quotes to enclose them.

Note: The command string should specify an AppleScript to launch. If it is only a file name, such as "**Database Script**", ADSM looks for this file in the Scripts folder inside the ADSM folder. Otherwise, the command string should contain the full path of the file, such as "**La Pomme:Scripting:Database Script**".

Examples

Preferences file example: postschedulecmd "restart database"

where the command string is a valid command for restarting your database.

PRESCHEDULECMD

PRESCHEDULECMD

Use the PRESCHEDULECMD option to specify a command that ADSM should process before running a schedule.

If you want ADSM to wait for the command to complete before continuing with other processing, specify PREschedulecmd. If you don't want it to wait, specify PRENschedulecmd.

Syntax

```
►► — [PREschedulecmd] — "cmdstring" ————— ►►  
      [PRENschedulecmd]
```

Parameters

"cmdstring"

The command to process.

You can use only one PRESCHEDULECMD option. You can also issue a command after a schedule using the POSTSCHEDULECMD option.

Use a blank or null string for *cmdstring* if you want to prevent any commands that the ADSM administrator uses for POSTSCHEDULECMD and PRESCHEDULECMD from running. If you specify a blank or null string on *either* option, it prevents the administrator from using a command on *both* options.

However, if your administrator uses a blank or null string on the PRESCHEDULECMD option, that prevents you from running a pre-schedule command.

If the command string contains blanks, you must enclose it in double quotes. If you have double quotes within the command string, use single quotes to enclose them.

Note: The command string should specify an AppleScript to launch. If it is only a file name, such as **"Database Script"**, then ADSM looks for this file in the Scripts folder inside the ADSM folder. Otherwise, the command string should contain the full path of the file, such as **"La Pomme:Scripting:Database Script"**.

Examples

Preferences file example: `preschedulecmd "quiesce database"`

where the command string is a valid command for quiescing your database.

QUERYSCHEDPERIOD

QUERYSCHEDPERIOD

Use the QUERYSCHEDPERIOD option to specify the number of hours you want the client scheduler to wait between attempts to contact the ADSM server for scheduled work.

This option applies only when the SCHEDMODE option is set to POLLING. ADSM uses this option only when the Scheduler is running.

Your ADSM administrator can also set this option. If your ADSM administrator specifies a value for this option, that value overrides what you specify in the Preferences file after your client node successfully contacts the ADSM server.

Syntax

►—QUERYSCHedperiod *hours*—►

Parameters

hours

The number of hours the client scheduler waits between attempts to contact the ADSM server for scheduled work.

The range of values is 1 to 9999.

The default is 12.

Examples

Preferences file example: QUERYSCH 6

QUIET

QUIET

Use the QUIET option to keep messages from being displayed on your screen in the Scheduler Status window.

This option only applies when you are running the Scheduler and ADSM is performing scheduled work.

When you use the QUIET option, certain error information still appears on your screen, and messages are written to log files.

If you do not specify QUIET, ADSM uses the VERBOSE option by default.

Syntax

►► Quiet ◄◄

Examples

Preferences file example: quiet

RETRYPERIOD

RETRYPERIOD

Use the RETRYPERIOD option to specify the number of minutes you want the client scheduler to wait between attempts to process a scheduled command that fails or between unsuccessful attempts to report results to the server.

ADSM uses this option only when the Scheduler is running.

Your ADSM administrator can also set this option. If your ADSM administrator specifies a value for this option, that value overrides what you specify in the Preferences file after your client node successfully contacts the ADSM server.

Syntax

▶—RETRYPeriod *minutes*—————▶

Parameters

minutes

The number of minutes the client scheduler waits between attempts to contact the ADSM server or attempts to process a scheduled command that fails.

The range of values is 1 to 9999.

The default is 20.

Examples

Preferences file example: RETRYP 10


SCHEDCOMPLETEACTION

SCHEDCOMPLETEACTION

Use the SCHEDCOMPLETEACTION to specify an action to take after a schedule has been completed.

ADSM uses this option only when the Scheduler is running.

Syntax



Parameters

Quit

Tells the scheduler application to quit once a schedule has completed.

SHUTdown

Causes your machine to be shut down once a schedule has completed. If there are other applications opened with unsaved data, these applications might prompt the user to save the data. This can prevent the machine from shutting down until the prompt is dismissed.

Examples

Preferences file example: `schedcompleteaction shutdown`

SCHEDLOGNAME

SCHEDLOGNAME

Use the SCHEDLOGNAME option to specify the name and location of a file where you want ADSM to store the schedule log.

ADSM uses this option only when the Scheduler is running.

Syntax

▶—SCHEDLOGname *filespec*—————▶

Parameters

filespec

The folder and file name where you want ADSM to store schedule log information when processing scheduled work. If you specify only a file name, ADSM stores the file in the default folder.

The default is the ADSM application folder and a file name of **ADSM Schedule Log**.

Examples

Preferences file example: schedlogname "SupDrive:David Sched Log"

SCHEDLOGRETENTION


SCHEDLOGRETENTION

Use the SCHEDLOGRETENTION option to specify:

- The number of days to keep entries in the schedule log
- Whether to save the pruned entries

The schedule log is pruned after a scheduled event completes.

Syntax



Parameters

N or *days*

How long to wait before pruning the schedule log.

N

Do not prune the log. This allows the log to grow indefinitely.

This is the default.

days

The number of days to keep log file entries before pruning the log.

The range of values is 0 to 9999.

D or **S**

Whether to save the pruned entries.

If you specify anything but D or S, ADSM uses the default D. You must use a space or comma to separate this parameter from the previous one.

D

Discard the log entries when pruning the log.

This is the default.

S

Save the log entries when pruning the log.

ADSM copies the entries pruned from the log to **ADSM Sched Pruned Log**.

Examples

Preferences file example: SCHEDLOGRETENTION 30 S

SLOWINCREMENTAL

SLOWINCREMENTAL

Use the SLOWINCREMENTAL option to specify a slower algorithm for processing incremental backups. With this option, ADSM backs up one directory at a time, and so uses less memory.

Use this option while performing incremental backups if your machine is running low on memory.

Syntax

▶—Slowincremental Yes No

Parameters

Yes

Your client node uses the method that uses less memory when processing incremental backups. You can also specify Yes as ON, True, or 1.

This is the default.

No

Your client node uses the fastest, most memory-intensive method when processing incremental backups. You can also specify No as OFF, False, or 0.

Examples

Preferences file example: slowincremental yes

TAPEPROMPT

TAPEPROMPT

Use the TAPEPROMPT option to specify whether you want ADSM to wait for a tape to be mounted if required for a backup, archive, restore, or retrieve process, or prompt you for your choice.

Syntax

▶▶—TAPEPrompt— No
 Yes

Parameters

No

ADSM does not prompt you for your choice. It waits for the appropriate tape to be mounted. You can also specify No as OFF, False, or 0.

This is the default.

Yes

ADSM prompts you when a tape is required for a backup, archive, restore, or retrieve process. At the prompt, you can choose whether to wait for the appropriate tape to be mounted. If you choose not to wait, the file is not processed. You can also specify Yes as ON, True, or 1.

Examples

Preferences file example: TAPEPROMPT NO

TCPBUFFSIZE

TCPBUFFSIZE

Use the TCPBUFFSIZE option to specify the size you want to use for the ADSM internal TCP/IP communication buffer.

A larger buffer can improve communication performance, but uses more memory.

Syntax

▶—TCPBuffsize *size*—————▶

Parameters

size

The size you want to use, in kilobytes, for the ADSM internal TCP/IP communication buffer.

The range of values is 1 to 32.

The default is 16.

Examples

Preferences file example: tcpb 2

TCPPOINT

TCPPOINT

Use the TCPPOINT option to specify a server's TCP/IP port address. This option is required in your Preferences file.

Obtain the port address from your ADSM administrator.

Syntax

▶—TCPPOINT *port_address*—————▶

Parameters

port_address

The TCP/IP port address used to communicate with an ADSM server.

The range of values is 1000 to 32767.

The default is 1500.

Examples

Preferences file example: tcpport 1501

TCPSERVERADDRESS

TCPSERVERADDRESS

Use the TCPSERVERADDRESS option to specify the TCP/IP address for an ADSM server. This option is required in your Preferences file.

Obtain this server address from your ADSM administrator.

Syntax

▶—TCPServeraddress *server_address*—————▶

Parameters

server_address

A 1 to 64 character TCP/IP address for an ADSM server.

The value you specify for this parameter can be a TCP/IP Internet domain name or a dot address.

Examples

Preferences file example:

```
tcps 129.33.10.429
tcps dsmchost.endicott.ibm.com
```

TCPWINDOWSIZE

TCPWINDOWSIZE

Use the TCPWINDOWSIZE option to specify the size of the TCP/IP sliding window you want to use for your client node. This is the size of the buffer used when sending or receiving data.

A larger size can improve communication performance, but uses more memory.

Syntax

►►—TCPWindowsize *window_size*—————▶▶

Parameters

window_size

The size you want to use, in kilobytes, for the TCP/IP sliding window for your client node.

The range of values is 1 to 128.

The default is 16.

Examples

Preferences file example: tcpwindowsize 1

TIMEFORMAT

TIMEFORMAT

Use the TIMEFORMAT option to select the format you want ADSM to use for displaying the system time.

Syntax

▶▶—TIMEformat *format_number*—————▶▶

Parameters

format_number

One of the following time formats:

- 1 Displays time in this format: 23:00:00
This is the default.
- 2 Displays time in this format: 23,00,00
- 3 Displays time in this format: 23.00.00
- 4 Displays time in this format: 12:00:00A/P

Examples

Preferences file example: TIMEFORMAT 4

VERBOSE

VERBOSE

Use the VERBOSE option to specify that you want processing information to appear on your screen in the Scheduler Status window.

This option only applies when you are running the Scheduler and ADSM is performing scheduled work.

The alternative is the QUIET option. If you do not specify either option, VERBOSE is the default.

Syntax

▶▶—Verbose—————▶▶

Examples

Preferences file example: VERBOSE

VERBOSE

The terms in this glossary are defined as they pertain to the ADSM library. If you do not find a term you are looking for, you can refer to the *IBM Dictionary of Computing*, McGraw-Hill, 1994. In the United States and Canada you can order this publication by calling McGraw Hill at 1-800-2MC-GRAW.

This glossary may include terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42nd Street, New York, New York 10036.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC2/SC1).

A

absolute. A copy group mode value that indicates that a file is considered for incremental backup even if it has not changed since the last time it was backed up. See *mode*. Contrast with *modified*.

active policy set. The policy set within a policy domain that contains the most recently activated policy. This policy set is used by all client nodes assigned to the current policy domain. See *policy set*.

active version. The most recent backup copy of a file stored in ADSM storage for a file that currently exists on a file server or workstation. An active version remains active and exempt from deletion until it is replaced by a new backup version, or ADSM detects during an incremental backup that the user has deleted the original file from a file server or workstation.

administrative client. A program that runs on a file server, workstation, or mainframe. This program allows administrators to control and monitor ADSM servers through ADSM administrator commands. Contrast with *backup-archive client*.

administrator. A user who has been registered to the server as an administrator. Administrators can be assigned one or more privilege classes. Administrators can use the administrative client to enter ADSM server commands and queries according to their privileges.

ADSM. ADSTAR Distributed Storage Manager

ADSTAR Distributed Storage Manager (ADSM). A client/server program that provides storage management and data access services to customers in a multivendor computer environment.

Advanced Program-to-Program Communications (APPC). An implementation of the SNA LU6.2 protocol that allows interconnected systems to communicate and share the processing of programs. See *Systems Network Architecture Logical Unit 6.2* and *Common Programming Interface Communications*.

APPC. Advanced program-to-program communications.

archive. A function that allows users to copy one or more files to a long-term storage device. Archive copies may be accompanied by descriptive information, may imply data compression software usage, and may be retrieved by archive date, file name, or description. Contrast with *retrieve*.

archive copy. A file or group of files residing in an archive storage pool in ADSM storage.

archive copy group. A policy object that contains attributes which control the generation, destination, and expiration of archive files. The archive copy group belongs to a management class.

archive retention grace period. The number of days ADSM retains an archive copy when the server is unable to rebind the file to an appropriate management class.

authentication. The process of checking and authorizing a user's password before allowing that user access to the ADSM server. Authentication can be turned on or off by an administrator with system privilege.

authorization rule. A specification that allows another user to either restore or retrieve a user's files from ADSM storage.

B

backup. A function that allows users to copy one or more files to a storage pool to protect against data loss. Contrast with *restore*.

backup-archive client. A program that runs on a file server, PC, or workstation and provides a means for ADSM users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

backup copy group. A policy object that contains attributes which control the generation, destination, and expiration of backup files. The backup copy group belongs to a management class.

backup retention grace period. The number of days ADSM retains a backup version when the server is unable to rebind the file to an appropriate management class.

backup version. A file, directory, or file space that a user has backed up that resides in a backup storage pool in ADSM storage. Though there may be more than one backup version of a file in ADSM storage, only one is considered the active version. See *active version* and *inactive version*.

bindery. A database that consists of three system files for a NetWare 3.11 or 3.12 server. The files contain user IDs and user restrictions.

binding. The process of associating a file with a management class name.

C

central scheduling. A function that allows an administrator to schedule backup and archive operations from a central location. The operations can be scheduled on a periodic basis or on an explicit date.

client. A program running on a file server, PC, workstation or terminal that requests services of another program called the server. There are two types of ADSM clients: administrative client and backup-archive client. See *administrative client* and *backup-archive client*.

Client Access/400. A software product that supports advanced program-to-program communications (APPC) in the DOS, OS/2, and Microsoft Windows environments and provides a set of end user services.

client domain. The set of drives, file systems, or volumes selected by a user for processing during a backup or archive operation.

client node. A file server or workstation on which the backup-archive client program has been installed that has been registered with the server.

client options file. A user-editable file that contains processing options to identify ADSM servers, select communication methods, specify backup, archive,

restore, and retrieve options, define scheduling options, and to choose formats for date, time, and numbers. Also called the **dsm.opt** file.

client polling. A client/server communication technique where the client node queries the server for scheduled work.

client/server. A communications network architecture in which one or more programs (clients) request computing or data services from another program (the server).

client system options file. A file, used on UNIX clients, that contains a number of processing options which identify the ADSM servers to be contacted for services. This file also specifies communications options, backup and archive processing options, and scheduling options. Also called the **dsm.sys** file. See *client user options file*.

client user options file. A user-editable file, used on UNIX clients, that contains processing options to identify the ADSM server to contact, to specify backup, archive, restore, and retrieve options, and to select formats for date, time, and numbers. Also called the **dsm.opt** file. See the *client systems options* file.

closed registration. A registration process in which an ADSM administrator must register workstations as client nodes with the server. Contrast with *open registration*.

command line interface. A type of user interface where commands are specified on the command line. Contrast with *graphical user interface*.

Common Programming Interface Communications (CPIC). A programming interface that allows program-to-program communication using SNA LU6.2. See *Systems Network Architecture Logical Unit 6.2*.

communication method. The method by which a client and server exchange information. For ADSM backup-archive clients, the method can be 3270, PWSCS, SNA LU6.2, or TCP/IP. See *Programmable Workstation Communication Service*, *Systems Network Architecture Logical Unit 6.2*, and *Transmission Control Protocol/Internet Protocol*.

communication protocol. A set of defined interfaces that allows computers to communicate with each other.

copy group. A policy object that contains attributes which control the generation, destination, and expiration of backup and archive files. There are two kinds of copy

groups: backup copy group and archive copy group. Copy groups belong to management classes. See *frequency*, *destination*, *mode*, *retention*, *serialization*, and *version*.

CPIC. Common Programming Interface Communications.

D

default management class. A management class assigned to a policy set that is used to govern backed up or archived files when a user does not explicitly associate a file with a specific management class through the include-exclude list.

destination. A copy group attribute that specifies the storage pool to which a file will be backed up or archived. At installation, ADSM provides two storage destinations named BACKUPPOOL and ARCHIVEPOOL.

domain. See *policy domain* or *client domain*.

dsm.opt file. See *client options file* and *client user options file*.

dsm.sys file. See *client system options file*.

dynamic. A copy group serialization value that specifies that ADSM accepts the first attempt to back up or archive an object, regardless of any changes made during backup or archive processing. See *serialization*. Contrast with *shared dynamic*, *shared static*, and *static*.

E

error log. A text file written on disk that contains ADSM processing error messages. These errors are detected and saved by the ADSM server.

exclude. The process of identifying files in an include-exclude list to prevent them from being backed up whenever a user or schedule issues an incremental or selective backup operation.

expiration. The process by which files are identified for deletion because their expiration date or retention period has passed. Backed up or archived files are marked for

deletion based on the criteria defined in the backup or archive copy group.

F

file server. A dedicated computer and its peripheral storage devices that are connected to a local area network that stores both programs and files that are shared by users on the network.

file space. A logical space on the ADSM server that contains a group of files. In ADSM, users can restore, retrieve, or delete file spaces from ADSM storage. On systems such as DOS or OS/2, a file space is a logical partition and is identified by a volume label. On UNIX systems, a file space is a logical space that contains a group of files backed up or archived from the same file system (or part of a file system defined with the VIRTUALMOUNTPOINT option in the client system options file).

frequency. A copy group attribute that specifies the minimum interval, in days, between incremental backups.

fuzzy backup. A backup version of a file that might not accurately reflect what is currently in the file because ADSM backed up the file while the file was being modified.

fuzzy copy. An archive copy of a file that might not accurately reflect what is currently in the file because ADSM archived the file while the file was being modified.

G

gigabyte (GB). (1) One billion (10⁹) bytes. (2) When referring to memory capacity, 1 073 741 824 in decimal notation.

graphical user interface (GUI). A type of user interface that takes advantage of a high-resolution monitor, includes a combination of graphics, the object-action paradigm, and the use of pointing devices, menu bars, overlapping windows, and icons. Contrast with *command line interface*. See *windowed interface*.

GUI. Graphical user interface.

H

hierarchical storage management client. A program that runs on a workstation or file server to provide space management services. It automatically migrates eligible files to ADSM storage to maintain specific levels of free space on local file systems, and automatically recalls migrated files when they are accessed. It also allows users to migrate and recall specific files.

hive. On Windows NT, a registry subkey that corresponds to a set of files in the `\system32\config` directory.

HSM. Hierarchical Storage Management

I

inactive version. A copy of a backup file in ADSM storage that either is not the most recent version or the corresponding original file has been deleted from the client file system. Inactive backup versions are eligible for expiration according to the management class assigned to the file.

include-exclude file. On UNIX clients, a file that contains statements which ADSM uses to determine whether to back up certain files and to determine the associated management classes to use for backup or archive. See *include-exclude list*.

include-exclude list. A list of INCLUDE and EXCLUDE options that include or exclude selected files for backup. An EXCLUDE option identifies files that should not be backed up. An INCLUDE option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services. The include-exclude list is defined either in the include-exclude file (for UNIX clients) or in the client options file.

incremental backup. A function that allows users to back up files or directories from a client domain that are not excluded in the include-exclude list and that meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class assigned to each file. Contrast with *selective backup*.

IPX/SPX. Internetwork Packet Exchange/Sequenced Packet Exchange.

L

Local Area Network (LAN). A variable-sized communications network placed in one location. It connects servers, PCs, workstations, a network operating system, access methods, and communications software and links.

M

management class. A policy object that is a named collection of copy groups. A management class is associated with a file to specify how the server should manage backup versions or archive copies of workstation files. See *binding* and *copy group*.

migration. The process of copying a file from a local file system to ADSM storage and replacing the file with a stub file on the local file system.

mode. A copy group attribute that specifies whether a backup file should be created for a file that was not modified since the last time the file was backed up. See *absolute* and *modified*.

modified. A backup copy group attribute that indicates that a file is considered for backup only if it has been changed since the last backup. A file is considered changed if the date, size, owner, or permissions have changed. See *absolute* and *mode*.

N

Named Pipe. A type of interprocess communication which allows message data streams to be passed between peer processes, such as between a client and a server.

NDS. NetWare Directory Services.

NETBIOS. Network Basic Input/Output System. An operating system interface for application programs used on IBM personal computers that are attached to the IBM Token-Ring Network.

NetWare Directory Services (NDS). A database used by Novell NetWare to organize resources for an entire network of NetWare servers. The database itself can be distributed across the network and accessed by all servers. The NDS is available in NetWare Version 4.0 and higher.

NetWare Loadable Module (NLM). Novell NetWare software that provides extended server capability. Support for various ADSM and NetWare platforms are examples of NLMs.

Networking Services/DOS (NS/DOS). A software product that supports advanced program-to-program communications (APPC) in the DOS and Microsoft Windows 3.1 environments. With NS/DOS, communications applications on your workstation can 'talk to' partner applications on other systems that support APPC.

NLM. NetWare Loadable Module.

node. See *client node*.

node name. A unique name used to identify a workstation, file server, or PC to the server.

NS/DOS. Networking Services/DOS.

NTFS. The NT file system on Windows NT.

O

open registration. A registration process in which users can register their own workstations or PCs as client nodes with the server. Contrast with *closed registration*.

owner. The owner of backup-archive files sent from a multi-user client node, such as AIX.

P

pattern-matching character. See *wildcard character*.

PC Support/400. A software product that supports advanced program-to-program communications (APPC) in the DOS, OS/2, and Microsoft Windows environments and provides a set of end user services.

policy domain. A policy object that contains one or more policy sets. Client nodes are associated with a policy domain. See *policy set*, *management class*, and *copy group*.

policy set. A policy object that contains a group of management class definitions that exist for a policy domain. At any one time, there can be many policy sets within a policy domain, but only one policy set can be active. See *active policy set* and *management class*.

Programmable Workstation Communication Services (PWSCS). A product that provides transparent high performance communications between programs running on workstations or on host systems.

PWSCS. Programmable Workstation Communication Services.

R

registration. The process of identifying a client node or administrator to the server by specifying a user ID, password, and contact information. For client nodes, a policy domain, compression status, and deletion privileges are also specified.

registry. On Windows NT, a central database that contains information about hardware, applications, and operating system settings for each machine on the network. Provides security and control over system, security, and account settings.

restore. A function that allows users to copy a version of a backup file from the storage pool to a workstation or file server. The backup copy in the storage pool is not affected. Contrast with *backup*.

retention. The amount of time, in days, that inactive backed up or archived files are retained in the storage pool before they are deleted. The following copy group attributes define retention: retain extra versions, retain only version, retain version.

retrieve. A function that allows users to copy an archive file from the storage pool to the workstation or file server. The archive copy in the storage pool is not affected. Contrast with *archive*.

S

scheduling mode. The type of scheduling operation for the server and client node. ADSM supports two scheduling modes: client-polling and server-prompted.

selective backup. A function that allows users to back up files from a client domain that are not excluded in the include-exclude list and that meet the requirement for serialization in the backup copy group of the management class assigned to each file. Contrast with *incremental backup*.

serialization. A copy group attribute that specifies whether a file can be modified during a backup or

archive operation and what to do if it is. See *static*, *dynamic*, *shared static*, and *shared dynamic*.

server. A program running on a mainframe, workstation, or file server that provides shared services such as back up and archive to other various (often remote) programs called clients.

server-prompted scheduling. A client/server communication technique where the server contacts the client node when work needs to be done.

session. A period of time in which a user can communicate with a server to perform backup, archive, restore, or retrieve requests.

shared dynamic. An ADSM copy group serialization mode. This mode specifies that if a file changes during backup or archive and continues to be changed after a number of retries, the last retry commits the file to the ADSM server whether or not it changed during backup or archive.

shared static. A copy group serialization value that specifies that a file must not be modified during a backup or archive operation. ADSM attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See *serialization*. Contrast with *dynamic*, *shared static*, and *static*.

SNA LU6.2. Systems Network Architecture Logical Unit 6.2.

space management attributes. Attributes contained in a management class that specify whether automatic migration is allowed for a file, whether selective migration is allowed for a file, how many days must elapse since a file was last accessed before it is eligible for automatic migration, whether a backup version of a file must exist before the file can be migrated, and the ADSM storage pool to which files are migrated.

static. A copy group serialization value that specifies that a file must not be modified during a backup or archive operation. If the file is in use during the first attempt, ADSM will not back up or archive the file. See *serialization*. Contrast with *dynamic*, *shared static*, and *static*.

storage pool. A named set of storage volumes that is used as the destination of backup or archive copies.

system drive or partition. On Windows NT, the drive or partition on which Windows NT is installed. The

Windows NT directory is usually called WINNT or WINNT35.

Systems Network Architecture Logical Unit 6.2 (SNA LU6.2). A set of rules for data to be transmitted in a network. Application programs communicate with each other using a layer of SNA called Advanced Program-to-Program Communication (APPC).

T

TCP/IP. Transmission Control Protocol/Internet Protocol

timeout. (1) An event that occurs at the end of a predetermined period of time that began at the occurrence of another specified event. (2) A time interval allotted for certain operations to occur; for example, response to polling or addressing before system operation is interrupted and must be restarted. (3) A terminal feature that logs off a user if an entry is not made within a specified period of time.

Transmission Control Protocol/Internet Protocol (TCP/IP). A standard set of communication protocols that supports peer-to-peer connectivity of functions for both local and wide-area networks.

V

version. The maximum number of different backup copies of files retained for files. The following backup copy group attributes define version criteria: versions data exists and versions data deleted.

W

wildcard character. An asterisk (*) or question mark (?) character that is used to search for various or certain combinations of alphanumeric and symbolic names. These names can reflect file names or character strings within a file, for example.

windowed interface. A type of user interface that is either a graphical user interface or a text-based interface. The text-based interface maintains a close affinity to the graphical user interface, including action bars and their associated pull-down menus and windows. See *graphical user interface*.

workstation. A programmable high level workstation (usually on a network) with its own processing hardware such as a high-performance personal computer. In a

local area network, a personal computer that acts as a single user or client. A workstation can also be used as a server.

Index

A

- absolute mode 54
- access privileges
 - file sharing 19
 - for AppleShare 18
- activation date/time 51
- active files
 - display 24
 - displaying 7
- active policy set 51
- active version 24
- ADSM
 - icon 5
 - understanding options 59
- ADSM Preferences file 59
 - See also* Preferences file
- ADSM Preferences-Sample file 59
- ADSTAR Distributed Storage Manager
 - ending 9
 - introducing ix, 1
 - online help 8
 - password 5, 8
 - publications x
 - using 5
- AppleShare information
 - backing up 18
- archive
 - advanced 39
 - compared to backup 11
 - deleting 39
 - deleting restore/retrieve volumes 33
 - description 37
 - file specification 35
 - mode 54
 - options window 36
 - retrieve 41
 - tree 38
- archive (A) attribute 17
- archive copy group
 - attributes, descriptions of 52
 - description of 51
- Archive Options window 36
- archiving files
 - specifying management class 69
 - summary of options 60

- assigning a management class
 - to folders 56
- attributes
 - archive copy group 52
 - backup copy group 52
- authorization
 - description 30
 - summary of options 63
- authorizing
 - user for restore or retrieve 30
- automatic backup and archive services
 - description of 3

B

- backing up
 - open file 19
- backing up files
 - specifying management class 69
 - summary of options 60
- backup
 - active version 24
 - AppleShare information 18
 - AppleShare system software 18
 - compared to archive 11
 - deleting restore/retrieve volumes 33
 - failures 15
 - file specification, by 20
 - folder, by 22
 - how managed 18
 - inactive version 24
 - include-exclude list 15
 - incremental
 - advanced 16
 - basic 12
 - partial 16
 - processing 17
 - types of 12
 - mode 54
 - new or changed files 12
 - overview 11
 - restore 24
 - selective
 - advanced 23
 - processing 24
 - status 14

- backup (*continued*)
 - volume label 16
- backup copy group
 - attributes, descriptions of 52
 - description of 50
- Backup Status window 14
- backup versions
 - fuzzy backup 54
- backup-archive client 1
- binding management classes to files 57
- BookManager 71
- books
 - BOOKS option 71
 - online access 63
- BOOKS option 71
- bottom up processing 67
- bytes transferred 14

C

- case sensitivity 19
- central scheduling
 - description of 3
 - setting up 45
 - summary of options 61
- changed file 17
- changed files, backing up 12
- changing
 - authorizations 30
 - sort order for lists of files 7
- CHANGINGRETRIES option 73
- client options
 - POSTSCHEDULECMD 92
 - PRESCHEDULECMD 93
- client scheduler
 - starting automatically 5
- client scheduler, starting 45
- command processing
 - summary of options 62
- command syntax xiii
- COMMMETHOD option 74
- communication
 - MacTCP 60
 - required software 60
- communication with server
 - summary of options 60
- COMPRESSALWAYS option 75
- COMPRESSION option 76
- Copy Destination attribute 54

- Copy Frequency attribute 52
- copy groups
 - description of 50
- Copy Mode attribute 54
- Copy Serialization attribute 53
- copy type 52

D

- date
 - incremental by 12
 - sorting by 24
- DATEFORMAT option 77
- default domain 17
- default management class 50, 51
- deleting
 - archives 39
 - authorizations 30
 - restore/retrieve volumes 33
- description
 - archive 37
 - management class 52
- desktop
 - rebuilding 34
- Destination folder 29
- directory
 - management classes 78
- DIRMC option 78
- disaster recovery 32
- disk recovery 32
- displaying
 - active and inactive files 7
 - inactive backups 24
 - management class 51, 52
 - messages on screen 107
 - messages, stopping 95
 - online help 8
 - policy information 51
- domain
 - specifying 79
- DOMAIN option 17, 79
- dynamic serialization 53
- dynamic, shared serialization 53
- DynaText 71

E

- elapsed time 15
- ending a graphical user interface session 9

- error log
 - pruning 81
- error processing
 - summary of options 63
- ERRORLOGNAME option 80
- ERRORLOGRETENTION option 81
- EXCLUDE option 82
 - groups of files 65
 - include-exclude list 65
 - processing 67
 - system files 65
 - wildcard characters 66

F

- failures
 - backup 15
- file
 - archive
 - deleting 39
 - file specification 35
 - tree 38
 - backing up open 19
 - changed 17
 - restoring
 - another node's 32
 - to another workstation 32
 - retrieve 41
 - retrieving
 - another node's 32
 - to another workstation 32
- file name
 - precautions 19
 - sorting by 24
- file sharing
 - access privileges 19
- file space
 - See also* restore/retrieve volumes
 - names 16
- file specification
 - backup 20
 - restore 25
 - wildcard characters 23
- files
 - backed up 14
 - examined 14
 - how managed 18
 - marked inactive 14
 - multiple Preferences files 64
 - overlapping 33

- files (*continued*)
 - sorting 24
- folder
 - backup 22
 - restore 28
 - sorting by 24
- full incremental
 - overview 12
- fuzzy backup 54

G

- graphical user interface
 - displaying online help 8
 - ending a session 9
 - starting a session 5
 - using a mouse 6

H

- help
 - BOOKS option 71
 - displaying online 8
 - online manuals 63
- hiding ADSM on the desktop 9

I

- icon, ADSM 5
- inactive files
 - display 24
 - displaying 7
- inactive version 24
- INCLUDE option
 - defined 83
 - groups of files 65
 - include-exclude list 65
 - processing 67
 - specifying management class 69
 - wildcard characters 66
- include-exclude list
 - backup, used during 15
- incremental
 - by date 12, 16
 - performance 12
 - types of 12
- incremental backup
 - process 17

L

- label, volume 16
- LANGUAGE option 85
- log
 - pruning errors 81
 - pruning schedule 99

M

- management class
 - backup 18
 - binding to files 57
 - choosing for files 55
 - choosing for folders 56
 - default 50, 51
 - directory, specify for 78
 - displaying 51, 52
 - name 52
 - overriding default 69
 - rebinding files to 57
 - specifying with INCLUDE option 69
- manuals
 - BOOKS option 71
 - online access 63
- MATCHALLCHAR option 86
- MATCHONECHAR option 87
- MAXCMDRETRIES option 88
- message
 - displaying on screen 107
 - stop displaying 95
- modification date
 - sorting by 24
- modified mode 54
- mouse, using 6
- multiple Preferences files 64

N

- name
 - copy groups 52
 - management class 52
- new files, backing up 12
- NODENAME option 32, 89
- NUMBERFORMAT option 90

O

- online help
 - displaying 8

open file, backing up 19

options

- BOOKS 71
 - CHANGINGRETRIES 73
 - COMMMETHOD 74
 - COMPRESSALWAYS 75
 - COMPRESSION 76
 - DATEFORMAT 77
 - DIRMC 78
 - DOMAIN 79
 - ERRORLOGNAME 80
 - ERRORLOGRETENTION 81
 - EXCLUDE 65, 82
 - INCLUDE 65, 83
 - LANGUAGE 85
 - MATCHALLCHAR 86
 - MATCHONECHAR 87
 - MAXCMDRETRIES 88
 - NODENAME 89
 - NUMBERFORMAT 90
 - PASSWORD 91
 - QUERYSCHEDPERIOD 94
 - QUIET 95
 - RETRYPERIOD 96
 - SCHEDCOMPLETEACTION 97
 - SCHEDLOGNAME 98
 - SCHEDLOGRETENTION 99
 - setting in a file 63
 - SLOWINCREMENTAL 100
 - TAPEPROMPT 101
 - TCPBUFSIZE 102
 - TCPPORT 103
 - TCPSERVERADDRESS 104
 - TCPWINDOWSIZE 105
 - TIMEFORMAT 106
 - understanding 59
 - VERBOSE 107
 - wildcard characters 66
- Options window
- archive 36
- overlapping icons 33

P

- partial incremental
 - overview 12
 - running, by date 16
- password
 - changing 8
 - Preferences file, setting in 91

- password (*continued*)
 - using 5
- PASSWORD option 91
- performance
 - incremental backup 12
- policy
 - displaying 51
 - domain 49, 51
 - set 50, 51
- port address
 - specifying 103
- POSTSCHEDULECMD option 92
- Preferences file
 - ADSM Preferences file 59
 - ADSM Preferences-Sample file 59
 - creating multiple 64
 - understanding options 59
- PRESCHEDULECMD option 93
- processing options
 - description of 3
- prompt
 - restore destination 26
- pruning
 - error log 81
 - schedule log 99
- publication translations xi
- publications
 - BOOKS option 71
 - online access 63
 - order numbers x
 - softcopy x

Q

- QUERYSCHEDPERIOD option 94
- QUIET option 95

R

- reading syntax diagrams xiii
- rebinding files to different management class 57
- rebuilding the desktop 34
- recover disk 32
- refreshing volume displays 8
- registering client 5
- restore
 - active version 24
 - advanced considerations 30
 - destination 26
 - file specification 25

- restore (*continued*)
 - folder 28
 - folder window 29
 - Get Info 34
 - inactive version 24
 - overview 24
 - parameters 26
 - rebuilding desktop 34
 - status 27
 - to following destination 26
 - to original location 26
 - Tree 27
 - wrong icons 33
- restore/retrieve volumes
 - delete 33
- restoring
 - another node's files 32
 - authorizing another user 30
 - workstation, to another 32
- Retain Extra Versions attribute 53
- Retain Only Version attribute 53
- Retain Versions attribute 54
- retention grace period 58
- retrieve
 - advanced 43
 - basic 41
- retrieving
 - another node's files 32
 - authorizing another user 30
 - workstation, to another 32
- retry backup 73
- RETRYPERIOD option 96

S

- SCHEDCOMPLETEACTION option 97
- SCHEDLOGNAME option 98
- SCHEDLOGRETENTION option 99
- schedule log
 - pruning 99
- scheduled services
 - displaying 46
 - running 45
 - setting up 45
- selective backup
 - file specification, by 20
 - folder, by 22
 - overview 20
 - process 24

- serialization
 - description of 53
- shared dynamic serialization 53
- shared static serialization 53
- SLOWINCREMENTAL option 100
- softcopy publications
 - order numbers x
- sorting
 - lists of files 7
- sorting lists 24
- Source folder 29
- starting
 - ADSM using different Preferences files 64
 - an ADSM GUI session 5
 - automatically 5
- starting a session
 - set password in Preferences file 91
- static serialization 53
- static, shared serialization 53
- status
 - backup 14
- stop displaying messages 95
- storage management policies 3, 49
- syntax diagrams xiii

T

- TAPEPROMPT option 101
- TCP/IP server address
 - specifying 104
- TCP/IP window size
 - specifying 105
- TCPBUFFSIZE option 102
- TCPPORT option 103
- TCPSERVERADDRESS option 104
- TCPWINDOWSIZE option 105
- TIMEFORMAT option 106
- trademarks vii
- transaction 15
- transfer rate 14
- translations, publication xi
- Tree
 - archiving 38
 - restore 27

U

- using the graphical user interface
 - changing sort order for lists of files 7
 - displaying active and inactive files 7

- using the graphical user interface (*continued*)
 - ending a session 9

V

- VERBOSE option 107
- Versions Data Deleted attribute 52
- Versions Data Exists attribute 52
- View menu
 - Refresh volumes option 8
- volume label
 - duplicate names 19
- volume name precautions 19
- volumes
 - refreshing displays 8

W

- wildcard characters 23, 65
 - examples of using 66
 - MATCHALLCHAR option 86
 - MATCHONECHAR option 87

Communicating Your Comments to IBM

ADSTAR Distributed Storage Manager

Using the Apple** Macintosh**
Backup-Archive Client
Version 2

Publication No. SH26-4051-00

If you especially like or dislike anything about this book, please use one of the methods listed below to send your comments to IBM. Whichever method you choose, make sure you send your name, address, and telephone number if you would like a reply.

Feel free to comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. However, the comments you send should pertain to only the information in this manual and the way in which the information is presented. To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

If you are mailing a readers' comment form (RCF) from a country other than the United States, you can give the RCF to the local IBM branch office or IBM representative for postage-paid mailing.

- If you prefer to send comments by mail, use the RCF at the back of this book.
- If you prefer to send comments by FAX, use this number:
 - United States: 1-800-426-6209
 - Other countries: (+1)+408+256-7896
- If you prefer to send comments electronically, use this network ID:
 - IBMLink from U.S. and IBM Network: STARPUBS at SJEVM5
 - IBMLink from Canada: STARPUBS at TORIBM
 - IBM Mail Exchange: USIB3VVD at IBMMAIL
 - Internet: starpubs@vnet.ibm.com

Make sure to include the following in your note:

- Title and publication number of this book
- Page number or topic to which your comment applies.

Readers' Comments — We'd Like to Hear from You

ADSTAR Distributed Storage Manager

Using the Apple Macintosh**
Backup-Archive Client
Version 2**

Publication No. SH26-4051-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.

Readers' Comments — We'd Like to Hear from You
SH26-4051-00



Cut or Fold
Along Line

Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
RCF Processing Department
G26/050
5600 Cottle Road
SAN JOSE, CA 95193-0001



Fold and Tape

Please do not staple

Fold and Tape

SH26-4051-00

Cut or Fold
Along Line



Program Number: 5648-020
5622-112
5697-078
24H3780
5763-SV1
5733-197
5686-073
5655-119
28H2250
28H2180
89G1342

Printed in U.S.A.

SH26-4051-00

